

Is Canada's Retirement Income System Working?¹

Jeremy Lise²

July 2001

¹I wish to thank Tim Sargent, Jeremy Rudin, Keith Horner, Steven James, Harriet Jackson, Tom Crossley, Paul-Henri Lapointe, and seminar participants at the Department of Finance for helpful discussions and comments. Any errors or omissions are my own. The views expressed in this paper are my own and should not be attributed to the Department of Finance.

²Economic Studies and Policy Analysis Division, Department of Finance, Canada.

Abstract

Using a pseudo-panel of cohort data and a straightforward decomposition of consumption, we address the question of whether Canadians are currently able to meet the goals set out in Canada's retirement income system, and whether they are likely to meet them in the future. Specifically, do Canadians enjoy a basic minimum standard of living in retirement, and are they able to avoid a serious disruption in livings standards at retirement? We find the answer is yes to both of these questions, for current and future retirees. The results are presented in a clear graphical form that facilitates the decomposition of consumption profiles into age, year and cohort components.

Résumé

À l'aide d'un pseudo-panel de données sur les cohortes et d'une simple décomposition de la consommation, nous abordons la question de savoir si les Canadiens sont actuellement en mesure d'atteindre les objectifs énoncés dans le système de revenu de retraite, et s'ils le seront à l'avenir. Plus précisément, les Canadiens ont-ils un niveau de vie minimal de base à la retraite et sont-ils en mesure d'éviter de subir une baisse sérieuse de leur niveau de vie à la retraite? Notre étude nous amène à répondre oui aux deux questions, et ce, dans le cas des retraites d'aujourd'hui et de demain. Les résultats sont présentés clairement dans une forme graphique, qui facilite la ventilation des profils de consommation par âge, année et composante de cohorte.

1 Introduction

There is anxiety in Canada among some that Canadians of the baby-boom generation will not be financially prepared for retirement over the next few decades. A quick scan of Canadian newspapers produces headlines such as “Reality check: They make \$200,000 a year, but aren’t saving enough for retirement” (Kane, *National Post* 1999), “Approaching the golden years – without the gold: Boomers courting financial disaster, warns finance guru” (Marck, *Edmonton Journal* 1999), and “Boomers face retirement bust: Sacrifices needed today: The middle-aged aren’t saving enough, experts worry” (Cornell, *National Post* 2000). Accompanying the press’s apparent concern about individual savings is the academic concern that national savings rates have fallen, and will continue to fall in the face of demographic changes over the next 30 years (Fougère and Mérette 1999). It should be noted that changes in aggregate national savings rates and individual rates are quite different matters, and will in general differ, even after the measurement differences have been dealt with.¹ The reason is that in aggregate the individual savings rates are weighted by income. Thus it is possible for the age profile of individual savings rates to remain constant, while aggregate rates change due to shifts in demographics, or redistribution of income from high to low savers.²

It is important to distinguish between national and individual savings rates, as they address different policy concerns. National savings is important in that it provides a good indicator for funds available for investment. At this level it does not matter if all these funds are provided by one saver or many savers. On the other hand, individual savings rates provide an indication of the future financial security of individuals. The low savings rate of one household is not offset in this case by the high savings rate of others. If we

¹The National Accounts measure of personal savings includes employer contributions to registered pension plans, changes in the net liability of life assurance companies, and savings by non-profit institutions—none of which are included in our measure of household savings.

²A thorough review of household savings, as well as the relationship to national savings, can be found in Browning and Lusardi (1996).

are concerned with the number of people who can provide for their retirement it is the individual savings rate that matters.

Whether or not Canadians are saving enough is an important question for public policy. The goal of Canada’s retirement income system, as expressed in the 1996 Federal Budget, is to provide “a basic income guarantee for seniors and [to help] Canadians avoid a serious disruption in their living standards at retirement.” An obvious concern is that if a large fraction of retirees from the baby-boom generation—who will be a large fraction of the total population—fail to adequately prepare for their retirement there may be substantial pressure to enrich the retirement income system. We address the degree to which Canadians currently meet the goals of basic minimum needs and maintenance of living standards using a straightforward decomposition of the consumption patterns of Canadians over the past 30 years.

Following the introduction, Section 2 justifies our use of consumption to define living standards and describes the data that this is constructed from. It then describes the process of constructing a pseudo panel of cohort data and the methodology of decomposing consumption profiles into age, cohort, and year components. Section 3 presents graphically the results of this consumption decomposition for the population as a whole, as well as by consumption quartile. As a check on adequacy of consumption, an indirect approach to addressing the concept of basic needs by examining the amount of consumption retirees devote to gifts is employed. Section 4 repeats the same decomposition for savings, to determine what can be expected for future retirees. Section 5 reconciles the anxiety suggested by this introduction with the optimistic results of the paper. Section 6 concludes.

2 Conceptual Framework

To answer the question: “Is Canada’s retirement income system working?” we start from the stated goals of the Canadian retirement income system. If individuals can meet a basic requirement and avoid a serious disruption of living standards at retirement we can consider it to be working. We turn next

to providing a framework for what should be considered a serious disruption in living standards.

2.1 Motivation

There is a close relationship between the goal of avoiding a substantial disruption in living standards and what economic theory tells us about the optimal age profile of consumption. As long as utility is time separable, a utility maximizing agent will want to equalize marginal utility across time periods. This does not imply that they will want to have equal consumption across all time periods, but for two adjacent periods it is going to be fairly close. The only reasons to deviate from equality would be the following: (a) Changes in household characteristics from one year to another (households will want to spend more when they have dependent children), (b) the household's discount rate is different from the interest rate (this is not likely to matter much over a short time period), (c) the household receives a surprise decrease in lifetime income (not very likely around retirement unless one is heavily invested in the stock market), or (d) the household has a rising income profile combined with a borrowing constraint (this will not be the case around retirement age).

With this relationship in mind we will be looking for a drop in the consumption profile—after adjusting for household size—around retirement, or a persistent decline in consumption during retirement. Such a break or decline would indicate an inability to maintain living standards through the retirement transition. We proceed in two steps to address whether or not Canadians are saving enough for retirement. In Section 3 we examine whether the currently retired experienced a drop in consumption at retirement, or a decline during retirement. Particular attention is given to potential differences due to different lifetime resources. Then in Section 4, we look at the behaviour of those not yet retired to infer whether they are more or less likely to experience a break in consumption at retirement, based on their savings behaviour over the past thirty years.

2.2 The Data

The data used in this analysis comes from the 1969, 1978, 1982, 1984, 1986, 1990, 1992, and 1996 Family Expenditure Surveys (FAMEX), and the 1997 and 1998 Surveys of Household Spending (SHS) conducted by Statistics Canada. Statistics Canada collects information on demographics, income, and detailed expenditures from a sample of Canadian households. The sample size varies between surveys from a low of 4,569 in 1984 to a high of 18,031 in 1969. In addition, the population of interest also varies by survey. Most of the surveys sample from both urban and rural households in the ten provinces, however, in 1984 and 1990 only households in the largest 15 census metropolitan areas (CMAs) were sampled. To maintain consistency we restrict the sample to households residing in major urban centres. A second adjustment is required due to a change in the survey unit of analysis. Prior to 1990 the unit of analysis was a “spending unit”—“a group of persons dependent on a common or pooled income for the major items of expense and living in the same dwelling or one financially independent individual living alone.” From 1990 onward the unit of analysis became the “household”—“a person or persons occupying one dwelling unit.” (Statistics Canada 1969, 1998) In order to maintain consistency we select only those spending units and households that comprise one economic family.³ The final restriction made to the sample is to retain only households in which the head is aged 25 to 75 inclusive.⁴ After making these selections we are left with observations on 63,071 urban households spanning the 30 years from 1969 to 1998.

2.3 Variables of Interest

As outlined above, we are primarily interested in the age profile of consumption and any trends in this by cohort. Additionally we consider saving, as

³For 1969 we are not provided with the number of economic families, so the selection is made by excluding unrelated individuals. This is not strictly comparable with the other survey years.

⁴In most of the surveys age 76 refers to age 76 and over, while age 24 refers to age 15 to 24.

it is one of the primary tools used by individuals to smooth consumption over the life-cycle. We also touch on disposable income for the purpose of illustrating the extra information we have available when using consumption. The definition of disposable income is straightforward. Disposable income is equal to market income plus government transfer payments less personal income taxes. Market income comprises wages and salaries, income from self-employment, investment income, miscellaneous income, and other money receipts.⁵

The definition of consumption is slightly more complicated. The economic definition will include all goods and services consumed in a defined period. For services and nondurable goods this is simply equal to expenditures on these items. For durable goods however, consumption is the flow of services provided by the stock over the period under study, in this case a year. Following Pendakur (1998) we define consumption as current consumption of nondurables plus a consumption flow from the stock of durables.⁶ Nondurable consumption is defined as total expenditure less expenditure on indirect savings and durable goods. From total expenditure we subtract all investment, life insurance payments, and public and private pension payments—as these reflect deferred consumption. Spending on shelter is removed as this is a major source of savings for many households. Finally, we subtract spending on automobile purchases, which provide consumption for many years.

Consumption is then calculated as nondurable consumption plus an imputed consumption flow for two major durable goods: housing and automobiles. An OLS regression of rent on indicators for survey year, location, number of rooms, and their interaction is estimated on the sample of renters. This equation is then used to estimate an imputed rent for home owners given year, location and number of rooms of their dwellings. The consumption flow from shelter is then equal to actual rent for renters, and imputed

⁵Other money receipts include receipts from money gifts, inheritances, lump sum settlements, lottery winnings etc.

⁶Semi-durables, such as clothing, are treated as nondurables here; we restrict durables to include only the two most important: housing and automobiles. For a thorough discussion of durables see Browning and Crossley (1999).

rent for home owners. A similar procedure is used to impute a consumption flow from automobile purchases for those who had auto expenses in excess of \$100 in the survey year. The automobile imputation is augmented to account for the difference both in probability of making a purchase and the value of the purchase conditional on household income and size.⁷

There are two possible ways to define savings in the FAMEX and SHS data: income minus consumption and net change in assets. Theoretically these should provide identical values for saving. Unfortunately, due to recall error, there can be substantial differences for individual households. Burbidge and Davies (1994) report that for the survey years 1978 to 1990 these differences are small in aggregate. As we will see in Section 2.4 it is the aggregate numbers that will allow us to describe the consumption and savings profiles for different birth cohorts. For the purpose of this analysis we present results for the net change in assets in the body of the paper and the corresponding results for income minus consumption in the appendix.

Income and consumption data is converted to real 1992 dollars using the consumer price index (CPI) for all expenditures.⁸

As mentioned above, the FAMEX and SHS are household surveys. What we are interested in, however, is the individual age profile of consumption. This requires a method of controlling for household size since the level of consumption available to an individual within a household will depend on how many other household members also have a claim on common resources. When making this adjustment we wish to take into account the economies of scale available to multi-member households. For example, the heating bill is not twice as much for two people as it is for one person occupying a similar apartment. To account for these economies of scale we can use an equivalency scale that converts household resources into the equivalent resources of a single individual.

⁷The definition and construction of these variables follows Pendakur (1998).

⁸This should be thought of as a first step, especially if we wish to interpret consumption as welfare. Given that we have detailed information on expenditures it is possible to construct a price index for each household that is weighted by the household expenditure shares.

Using total household consumption assumes full economies of scale (additional members require no additional expenditures) while dividing by the number of household members assumes no economies of scale (each additional member requires an equal share of expenditure). We use an equivalency scale that falls between these bounds—the square root of household size—to transform household consumption into equivalent individual consumption. This adult equivalent level of consumption is assigned to the household head, whose age is used in constructing the age profiles of consumption and income. Given the data selection outlined above and the use of the age of the household head, the decompositions will not strictly reflect the Canadian population at large. They should be thought of as representing the profiles for household heads residing in major Canadian urban centres.

2.4 Constructing Cohort Panel Data

The analysis of consumption dynamics requires observations on the same unit of analysis over time. Working at the individual level would require panel data. The FAMEX and SHS surveys are not panel surveys, and there is no way to follow the consumption patterns of individual households through the surveys. We can circumvent this problem by working at a slightly higher level of aggregation as suggested by Browning, Deaton, and Irish (1985) and elaborated on in Deaton (1985). By defining the unit of analysis to be a birth cohort—all those born in the same year—we can create panel data at the cohort level using a series of household cross-sectional surveys from different years.

While we are unable to analyse the patterns of individual households over time, we are able to examine the average behaviour of a group of households having a fixed membership. In the context of the expenditure surveys used here we can follow, for example, the average behaviour of those born in 1944 for the 30 years spanning 1969 to 1998. This provides us with the average behaviour of the 1944 cohort for ages 25, 34, 38, 40, 42, 46, 48, 52, 53, and 54 that we observe in the ten survey years from 1969 to 1989. Similarly, for those born in 1924, we observe their average behaviour as they move from age

45 in 1969 to age 74 in 1998. For each birth cohort we observe a different 30 year age span, with significant overlap between many cohorts. This overlap allows us to simultaneously identify a common age profile and the cohort trends.

Although the analysis presented here deals with the cohort averages, it would not be possible without the original micro level data. The reason for this is that by using the micro level data we are able to construct averages that are nonlinear in the data. For example we will make use of the average of log consumption rather than the log of average consumption. Similarly, when looking at savings rates, we want the average of the ratio of household savings to income, not the ratio of the averages. Access to the micro data also allows us to use the weighting factors in creating the means and allows us to weight the cohort cells by cohort size⁹ in the descriptive regressions that follow. From the data on the original 63,071 households included by our selection criteria we construct 494 cohort year observations. The cohort cell sizes for these observations are presented in Table 1.

2.5 Age, Cohort and Year Decomposition

The question that motivated this analysis is whether Canadians are able to avoid a substantial disruption in living standards around retirement. In section 2.4 we outlined the methodology of using cohort means to explore the dynamics of consumption around retirement age. The main reason we employ this technique is to separate the age profiles of consumption from cohort trends. In other words, we want to distinguish between a consumption profile that reflects age, and the possibility that cohorts born in different years will have higher (or lower) levels of lifetime consumption.

To make this distinction clear we begin by looking at the age profile of disposable income as an example, before turning to consumption. The log of disposable income is plotted against age of household head in Figure 1 below. The dotted line connects the points observed in 1969, while the dashed

⁹The cohort year means are weighted by the square root of cohort size to ensure larger




Figure 1: Cohort age profiles: Disposable income

line connects the points in 1998. These are the cross-sectional age profiles observed in the first and last survey years. The problem is that moving across age along either of these lines means we are looking at a sample based on different cohorts. In order to keep the reference group—the cohort—constant we look at the solid lines. These lines connect points corresponding to the same cohort, over the 30-year period from 1969 to 1998. These are the actual age profiles observed for the cohorts born in 1915, 1925, 1935, 1945, 1955, and 1965. Only every tenth cohort is presented for visual clarity. Representing the data in this way makes it easier to see the distinct age, birth cohort, and year effects present in the data.

We can immediately see several stylized facts about disposable income. Looking at the profiles for the cohorts we can see a similar hump-shaped age profile, rising to about age 50, and then declining. At the same time we see that the disposable income is higher at every age for those born in 1935 than it is for those born in 1925, and the same relationship holds between cohorts are given proper weight in the descriptive regressions.

the 1925 birth cohort and the 1915 birth cohort. Conversely, we see that at each age income is lower for those born in 1965 than it is for those born in 1955, which in turn is lower at each age than for those born in 1945. The cross-sectional age profiles for 1969 and 1998 are not able to illustrate these differences in lifetime income between cohorts. When we turn our attention to the consumption profiles this will be critical, as it is possible for the consumption profile for a cohort to be constant around retirement, but for the cross-sectional version to display declining consumption due to different lifetime income for different cohorts. Finally, we also observe a pattern of year to year fluctuations in each of the age profiles.

In order to simultaneously identify the age, cohort and year patterns in the data we are required to impose some structure on the relationships. The restrictions we impose are similar to Hall (1971) and closely follow Deaton (1997), where they are described in detail. In the decompositions that follow we assume that the variable of interest can be decomposed first into a typical age profile, second into secular trends that shift the entire position of age profiles for different birth cohorts, and third into year effects that temporarily cause all cohorts to deviate from their profiles. For a given cohort born in year c and observed in survey year t , we can write

$$y_{ct} = \beta + \alpha_a + \gamma_c + \psi_t + u_{ct}.$$

where y is the variable of interest (in this case disposable income) and the subscripts a , c , and t refer to age, cohort and, year respectively. This provides a straightforward way to decompose the data; however, it precludes the possibility of interactions between the three effects. To offset the restrictions imposed through additive separability we allow the data to choose the patterns for the age, cohort and year effects by using dummy variables for all three. This model can then be written in the form

$$\mathbf{y} = \beta + \mathbf{A}\alpha + \mathbf{C}\gamma + \mathbf{T}\psi + \mathbf{u}$$

where \mathbf{y} is the stacked vector of cohort observations through time, and \mathbf{A} , \mathbf{C} and \mathbf{T} are matrices of dummy variables for age, cohort and year respectively.

The obvious problem with trying to estimate this model is that there is an exact linear relationship between age, cohort and year (age plus cohort equals survey year) that carries over into an exact linear relationship in the dummy variables. The solution adopted here is to place restrictions on the year effects, making it possible to identify the three effects separately. We attribute all of the life cycle pattern to an age effect, all of the secular growth to changes in permanent income for different cohorts, and yearly macro economic fluctuations to the year effects. To do this we restrict the year effects to be cyclical; they are restricted to sum to zero over the period of analysis, and be orthogonal to a time trend.¹⁰

This decomposition is represented graphically in Figure 2 for disposable income. The upper left panel simply replicates Figure 1 without the cross-sectional lines. The lower left panel is the typical age profile, scaled to represent the cohort born in 1944. This age profile is then shifted up or down for different cohorts according to the cohort effect depicted in the upper right panel. The year effects, in the remaining panel, are orthogonal to a time trend and sum to zero as described above. The use of dummy variables as opposed to high-order polynomials results in profiles that retain some noise, but the overall shape comes out strongly. A high-order polynomial could likely produce the same profile, however, this would assume a continuous profile. When we look at consumption we are interested in potential discontinuities around retirement, a possibility that would be assumed away using a polynomial representation.

¹⁰While this transformation allows us to separate out statistically the age, cohort and year effects (effectively giving conditional means), it is important not to overstate the usefulness of the exercise. As noted by Heckman and Robb (1985), it is not possible to interpret the coefficients in such a regression in the usual manner. We cannot think of the age coefficient as the effect of changing age, holding year of birth and survey year constant as this is internally inconsistent. To make these kinds of inferences requires that additional information be used to identify the year effects, such as the use of data on macro economic variables. In the current setting however the decomposition provides us with a very straightforward and informative method of describing the consumption patterns of Canadian cohorts.




Figure 2: Age, year and cohort decomposition of disposable income

The stylized facts about disposable income over the past 30 years are clearly evident in the panels. The decomposition nicely separates out the hump shaped age profile for income, the growth in lifetime income for those born between 1915 and 1945 and the stagnation of this growth for those born after 1945, and finally, the cyclical deviations that display downturns around the recessions of the early 1980s and 1990s. The ability of the decomposition to confirm what we know about income trends over the past thirty years provides us with more confidence when we look at consumption.

3 Consumption Profiles of Birth Cohorts

Having outlined the methodology and looked at disposable income as an example in Section 2, we now employ this graphical technique to address the question of whether Canadians avoid a significant disruption in living standards around retirement. Evidence in support of a disruption would be




Figure 3: Imputed consumption per equivalent adult

a significant change in consumption around age 65, similar to the pattern we see in Figure 2 for household disposable income.

3.1 Consumption

Looking at the age profile for imputed consumption in Figure 3 (lower left panel) reveals a very flat consumption profile from age 50 to 75, with no obvious disruption around the typical retirement ages of 60 to 65. It is worth noting once again that we have allowed the data to pick out the age profile, we are not imposing any parametric form. The lack of any break in consumption indicates that on average the Canadian retirement income system is meeting the goal of maintaining living standards at retirement.

The very flat consumption profile after age 50 is in stark contrast to the pattern of falling household disposable income that begins around age 45 (Figure 2). The difference in the shapes of the profiles demonstrates that

Canadians are able to separate the path of individual consumption from the path of household income. The hump-shaped profile for household disposable income is heavily influenced by the same shape in market income. The consumption profile for an individual within the household is then flattened in several ways. The first is the fact that household composition is intimately linked to the age of the head of the household. Household size will typically follow a hump shaped profile as the composition increases from single, to couple, to couple with children, and then shrinks as first the children leave, and then eventually one of the partners dies. For this reason, a hump-shaped profile for household disposable income can be consistent with a constant level of consumption for the individual household members. Attanasio and Browning (1995), looking at working age households only, find that once household formation is accounted for the age profile of consumption is remarkably flat. Secondly, the household responds to the hump-shaped income profile by saving or borrowing to separate the anticipated path of income from the desired path of consumption. Finally, by purchasing large durable goods when income is high, the household receives a consumption flow when older that is not purchased out of current income.¹¹ The effects of household formation, saving during the high earning years, and the value of imputed rent from owned housing acts both to dampen consumption pre-retirement, and increase it post-retirement, resulting in a very flat profile.

3.2 Quartile Breakdown

The age profile of consumption shown in Figure 3 does not display any disruption to consumption around retirement. Thus, on average, Canadians—

¹¹Because we are adjusting for household size, our measure of the consumption flow imputed to housing will rise when household size falls. For example, when the children move out the parents' imputed housing consumption increases as a direct result. This increase can be interpreted in several ways. One possibility is that the parents now receive a higher flow of services from the house—perhaps they convert a bedroom into an office. Alternatively, we can think of this additional consumption as the value placed on remaining in the family home—the alternative being moving into a smaller (cheaper) house and substituting towards other types of consumption.

more precisely, urban household heads—maintain consistent consumption levels from age 50 to 75. It is worth asking whether this pattern holds more broadly than just at the mean. Specifically we are concerned with the ability of those with lower lifetime resources to maintain their living standards through the retirement transition. To address this question, we reconstruct the cohort data, creating four subsets based on consumption quartile.

Households are assigned to quartiles as follows. For each survey year, the equivalent adult imputed consumption for a household is ranked against that of households in which the head is of the same age. In other words, we are ranking consumption within cohorts, by year. For example, in the 1990 survey year, for a household with head born in 1950, the consumption quartile is determined relative to other households in 1990 with a head born in 1950. As noted in Section 2.4, the cohort group must have a fixed membership over time. In order for the cohort quartiles constructed here to satisfy this fixed population assumption it must be the case that households maintain their within-cohort consumption rank over time. In other words, a household with consumption in the second quartile in 1969 will also be in the second quartile in 1998. This would clearly be too strong an assumption for income quartile, as it would assume away income dynamics. However, it is theoretically defensible for consumption, since according to the permanent income hypothesis, consumption in any given year is a function solely of lifetime wealth. We can appeal to the data for additional support for this assumption: Figures 3 and 2 illustrate that households are able to separate the path of consumption from the path of disposable income—at least after age 50.¹²

¹²Another strong candidate for subdividing the cohorts would be education level of the head, which is certainly fixed between the ages of 25 and 75. We present the corresponding results using education in the appendix, where the four education categories used are less than nine years of school, some high school, post-secondary diploma, and university degree. There are two reasons we prefer consumption quartile over education for inclusion in the body of the paper. Consumption quartile by definition provides an equal number of observations per quartile. Education groups do not divide the population evenly, and the composition changes substantially when we move from older to younger cohorts. The result is that the estimates obtained using education are much noisier. The second problem




Figure 4: Imputed consumption per equivalent adult by consumption quartile




Figure 5: Household disposable income by consumption quartile

Figures 4 and 5 repeat the decompositions of equivalent consumption and household disposable income by consumption quartile. The quartile decomposition of consumption is presented in Figure 4. The age and cohort effects for quartiles two, three and four basically mimic the population as a whole, while the first quartile is different slightly. For the first quartile, consumption appears to be quite consistent over the entire age profile. Comparing the consumption age profiles to the corresponding disposable income profiles in Figure 11 is suggestive of a slight disruption to the consumption path around retirement age, although to a much dampened extent relative to the disruption we see in disposable income.¹³ The quartile pattern for household disposable income displayed in Figure 5 is also interesting. Quartiles two, three and four replicate what we see for the population as a whole. The first quartile, however, is characterized by a discontinuity around age 65, at which time disposable income suddenly increases. This illustrates two points. First, that government retirement income programmes make up a substantial portion of household disposable income for the lowest consumption quartile.

is that education is not included as a variable on the public use files of the SHS for 1997 and 1998, reducing the number of usable years in the sample from ten to eight, and the number of cohort observations from 494 to 392. This means we need to estimate two fewer parameters, but lose 102 observations. That said, the results obtained by using education do not differ qualitatively from those obtained using consumption quartile, and are included as Figures 14 to 17 in the appendix.

¹³Plotting confidence intervals around the age profile indicates that we cannot reject the hypothesis that consumption is constant between ages 50 and 75 for each consumption quartile. The confidence bands have been omitted to simplify the graphs. Ideally we would construct a formal hypothesis test. Unfortunately a formal test based on the public use data employed in this analysis would not yield meaningful results as we cannot construct appropriate measures for the precision of the estimated coefficients due to the nonrandom sampling framework employed by Statistics Canada. That said, our null hypothesis is of equality of consumption between ages 50 and 75. We cannot reject this, even with the standard errors that are too small; using the correct (larger) standard errors then would lead to the same failure to reject. Constructing meaningful statistical tests requires the use of the master data files. For a detailed discussion of the issues involved in inference in the face of a nonrandom sampling framework see Deaton (1997) or chapter 30 of the Stata User's Guide (1999).

Second, household income in this quartile is actually higher post-retirement than in the years leading up to retirement.

The difference in the shapes of the cohort effects between the quartiles is also informative. For the population as a whole we saw a slowdown of growth in lifetime income for cohorts born after 1945. The breakdown into quartiles makes it clear that this stagnation for the post 1945 cohorts is the aggregate result of falling disposable income for the first quartile, no change for the second quartile, and minor growth enjoyed by the third and fourth quartiles.

What is clear from this graphical presentation, however, is that on average Canadians do not experience a disruption in living standards around retirement (Figure 3) and if the lowest quartile experiences a consumption shock around retirement it is likely positive rather than negative (Figure 4).

3.3 Gifts as a Check on Consumption Adequacy

In Section 3 we explored one goal of Canada’s retirement income system, that of helping Canadians avoid a significant disruption to living standards around retirement. We conclude that Canadians can indeed avoid such a disruption. We turn now to assessing whether the goal of a minimum living standard is met. Rather than defining either an absolute or relative basic living standard, we take an indirect approach that attempts to determine whether individuals behave as if their minimum needs are met. To do this we consider the share of consumption that is devoted to giving gifts to those outside the household. The motivation here is that for households in which basic household needs are not met, we will not expect to observe gifts to individuals or organizations outside the household.¹⁴

¹⁴This is not to suggest that gifts should be considered fundamentally different from other forms of consumption; gifts may be payment for services such as visits from family members or membership in social organizations. We make use here of the cross-sectional observation that the expenditure share devoted to gifts increases with income, indicating that gifts are a luxury good. Observing an increase in the expenditure share devoted to gifts would be inconsistent with a household not meeting its basic needs.




Figure 6: Share of household consumption devoted to gifts by quartile

Figure 6 plots the age profile for the average share of household consumption devoted to gifts, by consumption quartile. This profile is constructed in the same manner as before, controlling for year and cohort effects, with only the age profiles presented for brevity. The well known pattern emerges: gifts increase with age and wealth. While the first quartile displays lower giving rates than the others at all ages, the interesting comparison is within the quartile. There is a noticeable increase in the gift share after age 65, indicating that the transition to retirement does not constrain consumption for those with the lowest lifetime resources; there is no evidence of a pervasive struggle to preserve the level of consumption at retirement.¹⁵

¹⁵This exercise is repeated in Figure 19 in the Appendix using education group. In this case the difference between education groups is dampened relative to the difference between consumption quartiles. The patterns within education groups are qualitatively the same as within the consumption quartiles.

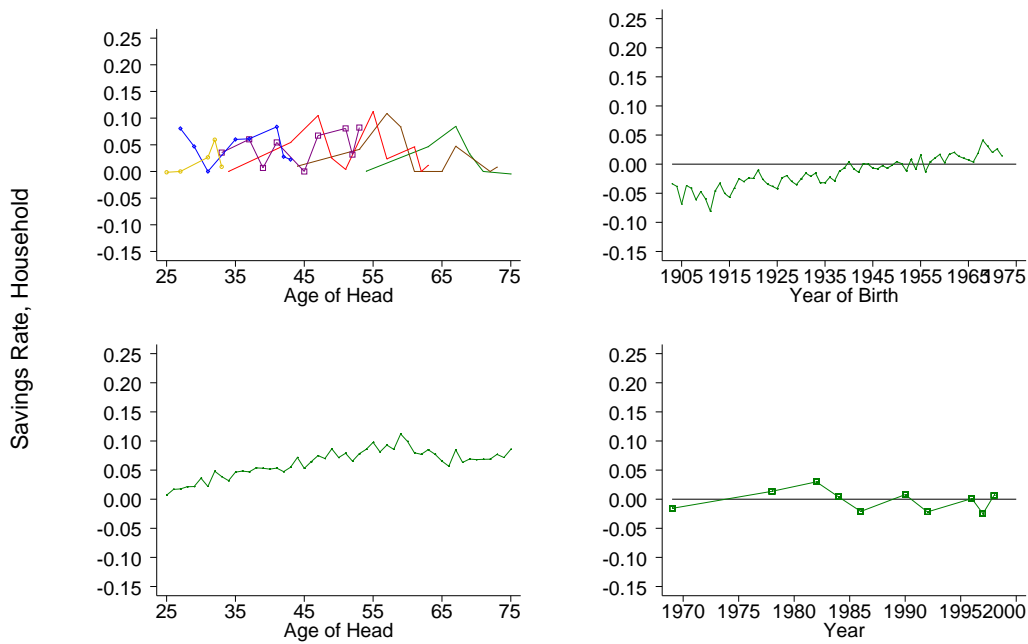


Figure 7: Household savings rates (net change in assets)

4 Savings and Future Retirees

4.1 Savings

The flat portion of the age profile of consumption in Figure 3 (bottom left graph) between ages 50 and 75 is identified using data on those cohorts that we observe in this age range, specifically those born between 1900 and 1945. From this we can infer that on average Canadians who have already retired have not experienced a disruption in their living standards. One of the motivating factors for this paper was the concern that the baby-boom may not be financially prepared for retirement, thus placing additional pressure on public programs when this large cohort retires. To address this concern we need to further examine the behaviour of those born after 1945.

Given the existing system of government income support programmes, the only reason for a break down would be if private savings were inadequate. From the results presented above we can conclude that retirement

savings have been adequate in the past to maintain retirement consumption. As long as those cohorts not yet retired have savings rates at least as high, maintaining consumption in retirement should not be a problem. Using the same approach that was applied to consumption, we can estimate a typical age profile of saving, with a level effect for each birth-cohort, and allowing for business cycle effects. In this case we are most interested in the cohort effect. Figure 7 decomposes the median household savings rate into age, cohort and year effects. The savings rate here is defined as the net change in assets divided by disposable income. The most noticeable difference between the savings decomposition versus the income or consumption decompositions is the much larger year to year variation. This is to be expected as households use savings as the tool to separate the anticipated paths of income and consumption. In the face of income shocks, savings can fluctuate to insulate consumption.

We can see in Figure 7 that the cohort effect is increasing in year of birth; those not yet retired have higher savings rates at every age than did the currently retired cohorts.^{16 17} The fact that younger cohorts have higher average savings rates indicates that as long as the return on retirement savings is at least as high as it is for the currently retired, private savings will be more than adequate for future retirees.

¹⁶Figure 12 in the Appendix uses the income minus consumption definition of savings with no qualitative difference.

¹⁷In the age profile of savings estimated here we do not appear to observe any dissavings by the elderly. This is partly due to a difficulty in defining savings rates using household survey data. The problem is that the definition of income used does not make a distinction between earned and other income. The data does not let us distinguish between consuming interest income and drawing down principle. For example, annuity income includes some interest income, but is largely principal. This issue is discussed at length in Fougère and Mérette (1999). For the purposes of this paper this is not really a problem. What we need to know about savings is whether there are differences between cohorts in the position of the age profile of savings, not what the actual profile looks like. Additionally, because we use dummy variables for age, the estimates for the older ages do not influence the estimates for younger ages, as would be the case with a polynomial.

4.2 Quartile Breakdown

Figure 8 plots the savings decompositions by consumption quartile. The first quartile is once again of interest, as its members display a flat age profile centred just above zero, and a flat cohort trend. In spite of this, this group does not suffer a deterioration in standard of living in old age. The story here is that those in the lowest quartile do not save at any age, and never have. The other three quartiles have age profiles reflecting positive savings. Most importantly, the cohort effect is increasing, suggesting that those currently in the workforce are saving more than their predecessors.

4.3 Smooth Consumption, Very Different Savings

In the previous section we saw that the lowest consumption quartile does not experience a decline in consumption in old age, despite the fact that they do not save. Additionally, the other three quartiles all have consumption profiles that share a common shape—albeit at different levels of consumption—but have quite different savings profiles. The apparent inconsistency can be reconciled by looking at the composition of income sources across the quartiles, particularly government transfers as a share of total income. Figure 9 plots the estimated age profiles of the share of total income received from government transfers—including Canada and Quebec Pension Plans (C/QPP)—by quartile.¹⁸ Focussing on ages 65 plus we see that, as expected, the importance of government transfers as a source of income varies by consumption quartile. Government transfers constitute approximately 80 per cent of total income for the lowest quartile, this falls to 30 per cent for the fourth quartile. The clear message is that at higher levels of lifetime consumption a much larger share of retirement income must come from non-transfer sources. The key point here is that for the lowest quartile there is no need to save as public retirement income support programs ensure consumption is fully replaced at

¹⁸The age profiles are estimated controlling for year and cohort effects in the same manner described earlier; we present only the age profile for brevity. The same exercise is repeated using education group in Figure 18 in the Appendix with identical results.




Figure 8: Household savings rate by consumption quartile




Figure 9: Share of household income from government sources by quartile

the pre-retirement level. For those with higher lifetime resources, the transfers portion of the retirement income system—including C/QPP receipts—is much less important than private sources for maintaining living standards through the transition to retirement.

5 Reconciling the Anxiety

The results in this paper are substantially more optimistic than the headlines quoted in the introduction. Additionally we are still left with the puzzle of falling private savings in the National Accounts. Is it possible to reconcile these results with the anxiety that motivated the work? Looking first at the issue of the National Accounts, we can see that it is actually quite easy to reconcile low personal savings with adequate household retirement savings. The important distinction is that the savings rate in the National Accounts is income weighted, while average household savings is population weighted

(the former is a ratio of the averages, while the latter is an average of the ratios). For this reason alone the average household savings rate will not in general equal the national personal savings rate. This is not to say that the national savings rate is not important, it certainly provides an indication of resources available for investment, but it cannot tell us whether on average households are well prepared for their retirement.¹⁹

When reading the headlines in the introduction it helps to identify for whom private savings is important for maintaining consumption in retirement. It is tempting to reason that if the household earning \$200,000 per year is not saving enough to retire, then the household earning \$20,000 and not saving at all will be in big trouble. This is not really true. What constitutes an adequate savings rate is not independent of lifetime resources. Government transfers make up approximately 80 per cent of total retirement income for the first consumption quartile, but only 25 per cent for the highest quartile. This means that only 20 per cent of retirement income comes from private sources for the lowest quartile. Since private sources include employer-sponsored pension plans as well as personal savings, savings plays a relatively small role in the retirement income for this group. In contrast, the highest quartile relies on private sources for 75 per cent of retirement income, suggesting a substantially greater role for personal savings.

6 Conclusion

In conclusion, it appears that the goals of the Canadian retirement income system are being met, and that the concerns expressed do not bear out in the data. Once household size and the imputed consumption value of owned housing is taken into account, consumption remains remarkably stable across the retirement transition and during retirement. This holds true regardless of consumption quartile. Additionally, households in the lowest quartile appear

¹⁹Lusardi, Skinner, and Venti (2001) emphasize that the National Accounts saving rate is not useful in judging whether households are preparing for retirement or other contingencies.

to have enough resources to meet basic needs. Indeed, a substantial fraction have enough to make gifts. Comparing the patterns of saving we observe for those who will retire in the next two decades relative to the savings we observed from the currently retired suggests that future retirees will also be able to avoid a disruption to consumption through the retirement transition.

References

- ATTANASIO, O., AND M. BROWNING (1995): “Consumption over the Life Cycle and over the Business Cycle,” *The American Economic Review*, 85(5), 1118–37.
- BROWNING, M., AND T. CROSSLEY (1999): “Shocks, Stocks and Socks: Consumption Smoothing and the Replacement of Durables During an Unemployment Spell,” Working Paper in Economics and Econometrics 376, Australian National University.
- BROWNING, M., A. DEATON, AND M. IRISH (1985): “A Profitable Approach to Labour Supply and Commodity Demands Over the Life-cycle,” *Econometrica*, 53(3), 503–43.
- BROWNING, M., AND A. LUSARDI (1996): “Household Saving: Micro Theories and Micro Facts,” *Journal of Economic Literature*, 34, 1797–1855.
- BURBIDGE, J., AND J. DAVIES (1994): “Household Data on Savings Behavior in Canada,” in *International Comparisons of Household Savings*, ed. by J. Poterba, chap. 1, pp. 11–56. University of Chicago Press.
- CORNELL, C. (2000): “Reality check: They make \$200,000 a year, but aren’t saving enough for retirement,” *National Post*, p. 21.
- DEATON, A. (1985): “Panel Data from Time Series of Cross-sections,” *Journal of Econometrics*, 30, 109–26.
- (1997): *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. The Johns Hopkins University Press, Baltimore.
- FOUGÈRE, M., AND M. MÉRETTE (1999): “An Econometric Examination of the Impact of Population Ageing on Personal Savings in Canada,” Working Paper 99–03, Department of Finance, Canada.

- Government of Canada (1996): “Budget 1996: The Seniors Benefit, Securing the Future,” Budget, Government of Canada.
- HALL, R. (1971): “The Measurement of Quality Change from Vintage Price Data,” in *Price Indexes and Quality Change*, ed. by Z. Griliches, chap. 8, pp. 240–71. Harvard University Press.
- HECKMAN, J., AND R. ROBB (1985): “Using Longitudinal Data to Estimate Age, Period and Cohort Effects in Earnings Equations,” in *Cohort Analysis in Social Research: Beyond the Identification Problem*, ed. by W. Mason, and S. Fienberg, chap. 5, pp. 137–50. Springer-Verlag, New York.
- KANE, M. (1999): “Boomers face retirement bust: Sacrifices needed today: The middle-aged aren’t saving enough, experts worry,” *National Post*, p. C6.
- LUSARDI, A., J. SKINNER, AND S. VENTI (2001): “Saving Puzzles and Saving Policies in the United States,” Working Paper 8237, NBER.
- MARCK, P. (1999): “Approaching the golden years – without the gold: Boomers courting financial disaster, warns finance guru,” *Edmonton Journal*, p. E2.
- STATA CORP (1999): *Stata User’s Guide Release 6.0* College Station, Texas.
- STATISTICS CANADA (1969, 1978, 1982, 1984, 1986, 1990, 1992 & 1996): *Survey of Family Expenditures: Public-Use Microdata File Users’ Guide*.
- (1997 & 1998): *Survey of Household Spending: Public-Use Microdata File Users’ Guide*.
- PENDAKUR, K. (1998): “Changes in Canadian Family Income and Family Consumption Inequality Between 1978 and 1992,” *Review of Income and Wealth*, 44(2), 259–83.

A Appendix




Figure 10: Disposable income per equivalent adult

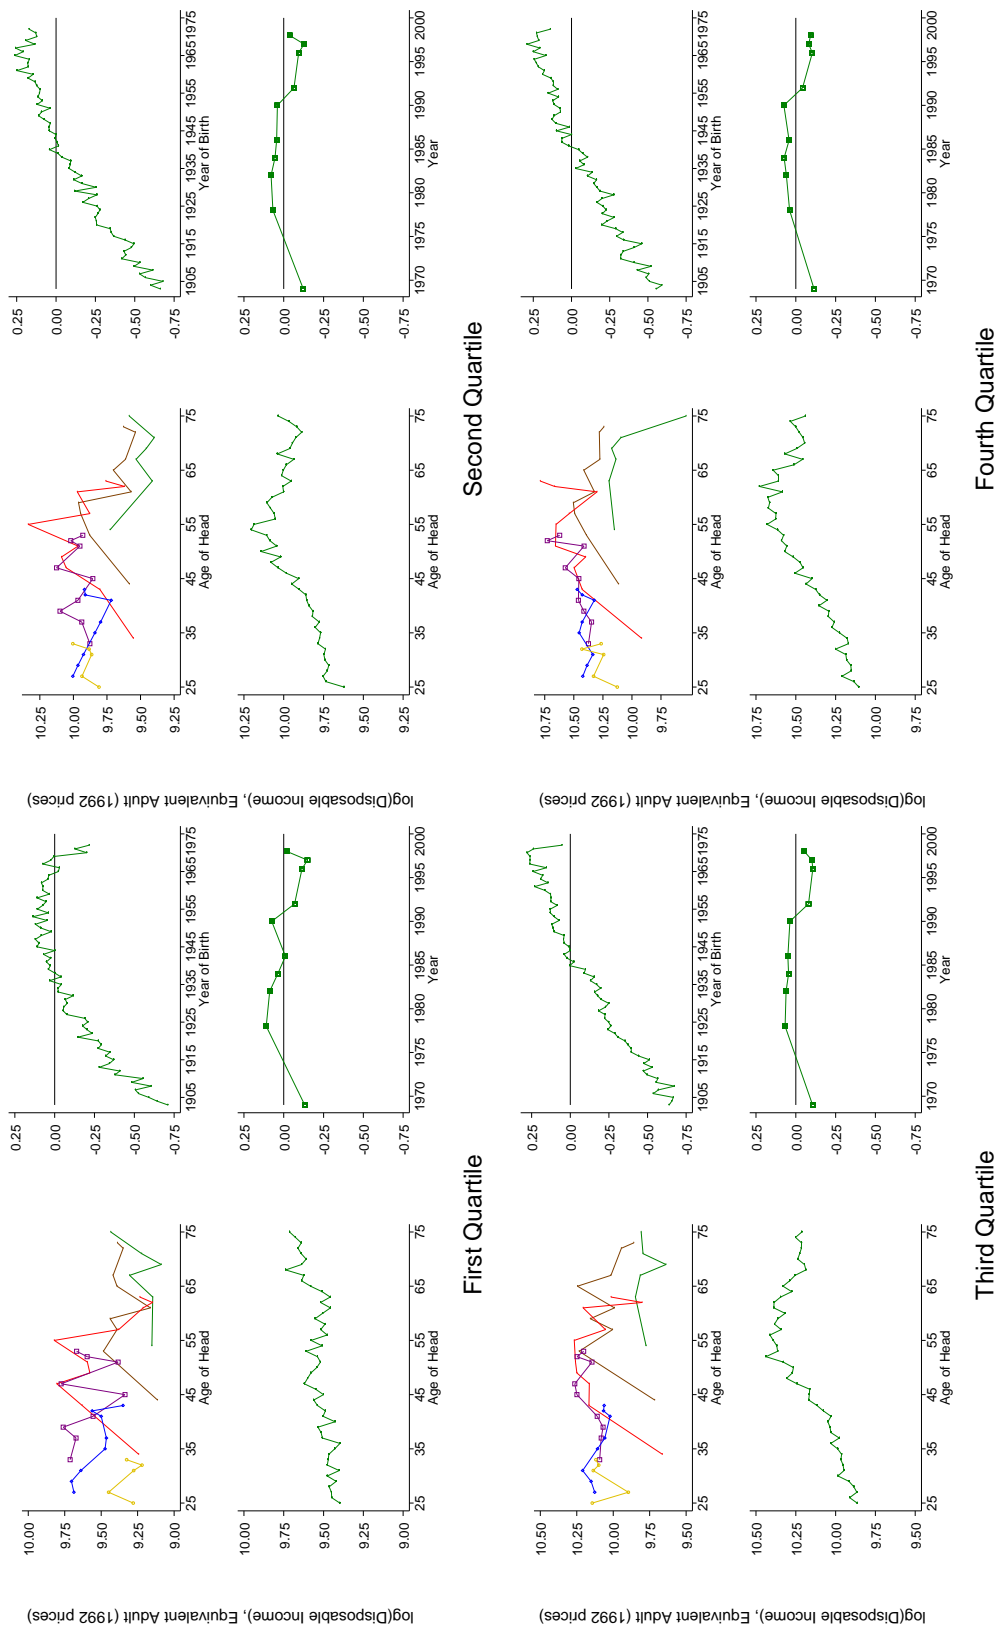


Figure 11: Disposable income per equivalent adult by consumption quartile

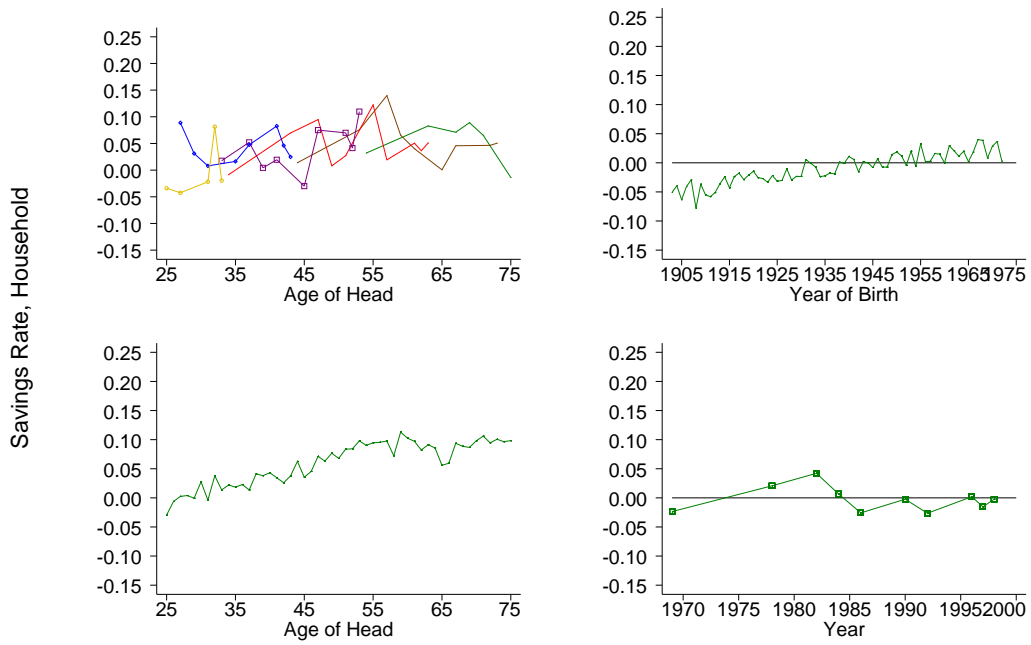


Figure 12: Household savings rate (income minus consumption definition)

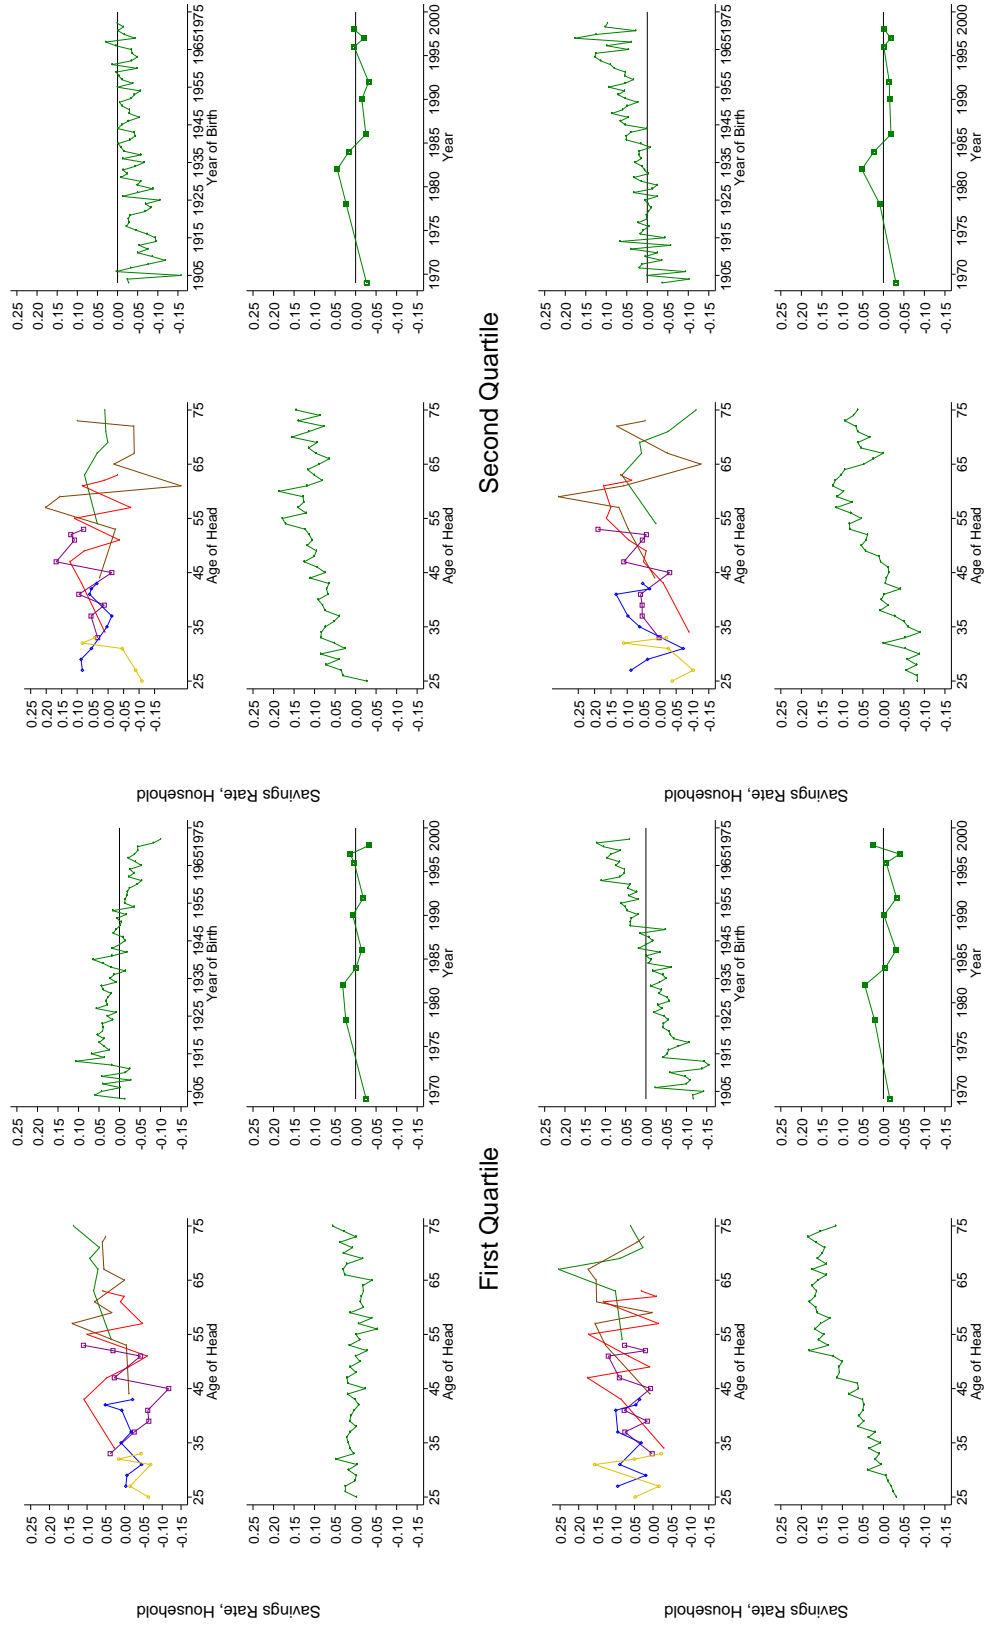


Figure 13: Household savings rate by consumption quartile (income minus consumption definition)




Figure 14: Household disposable income by education group

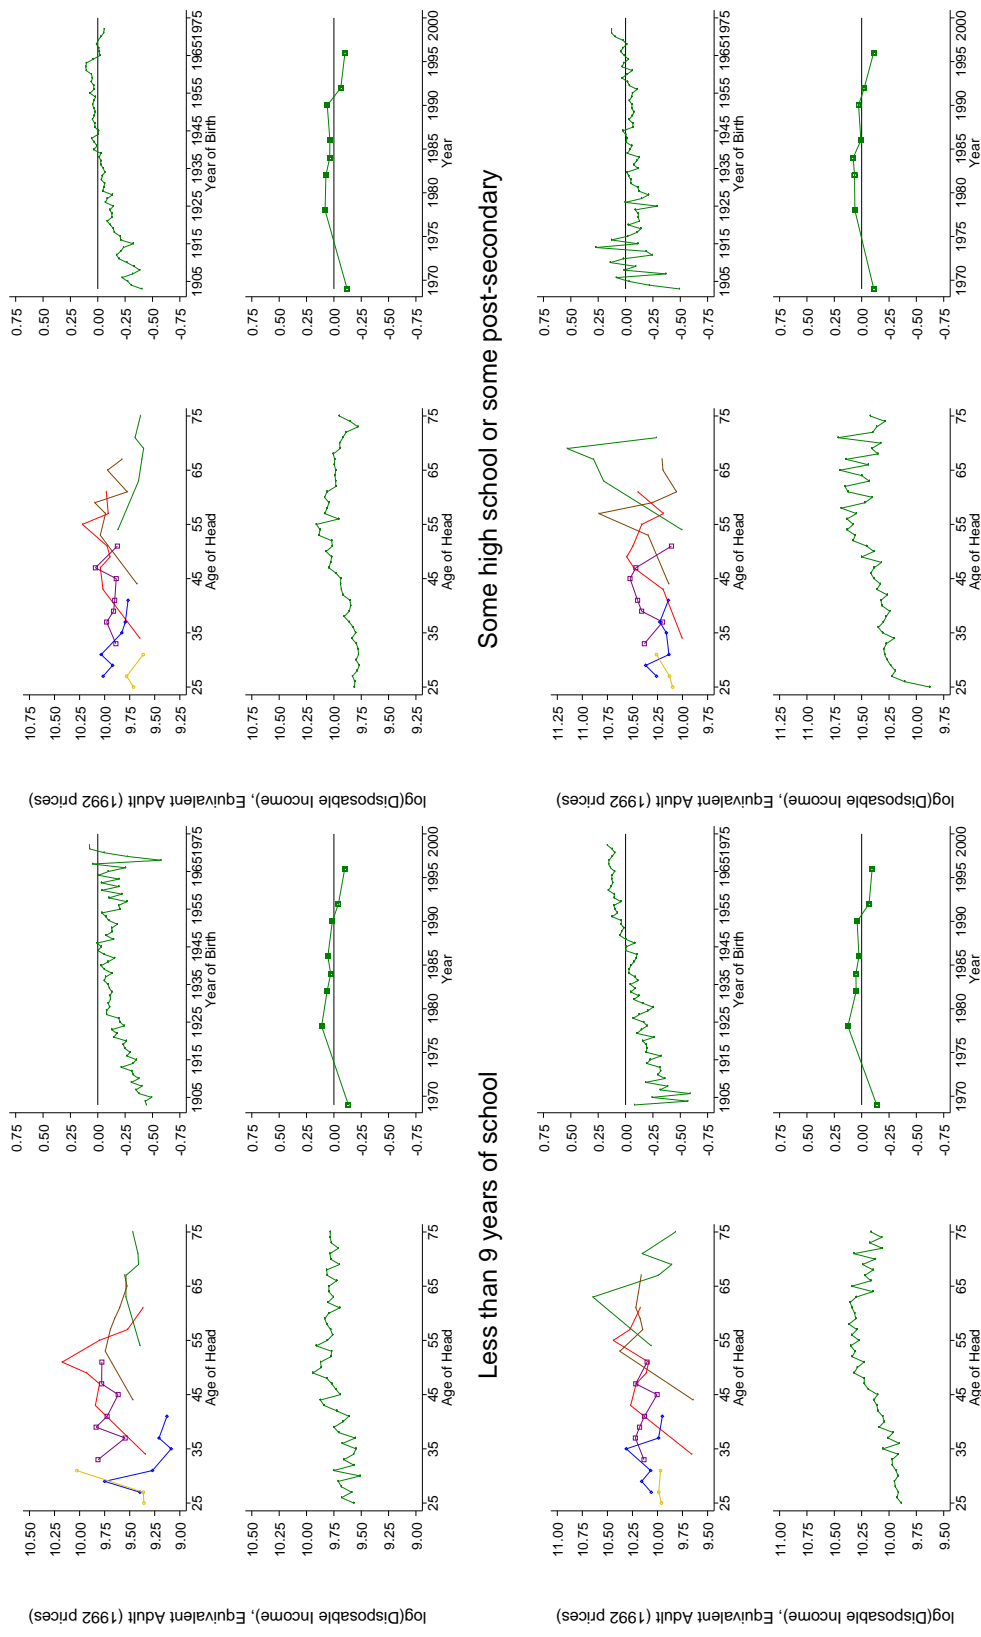


Figure 15: Disposable income per equivalent adult by education group

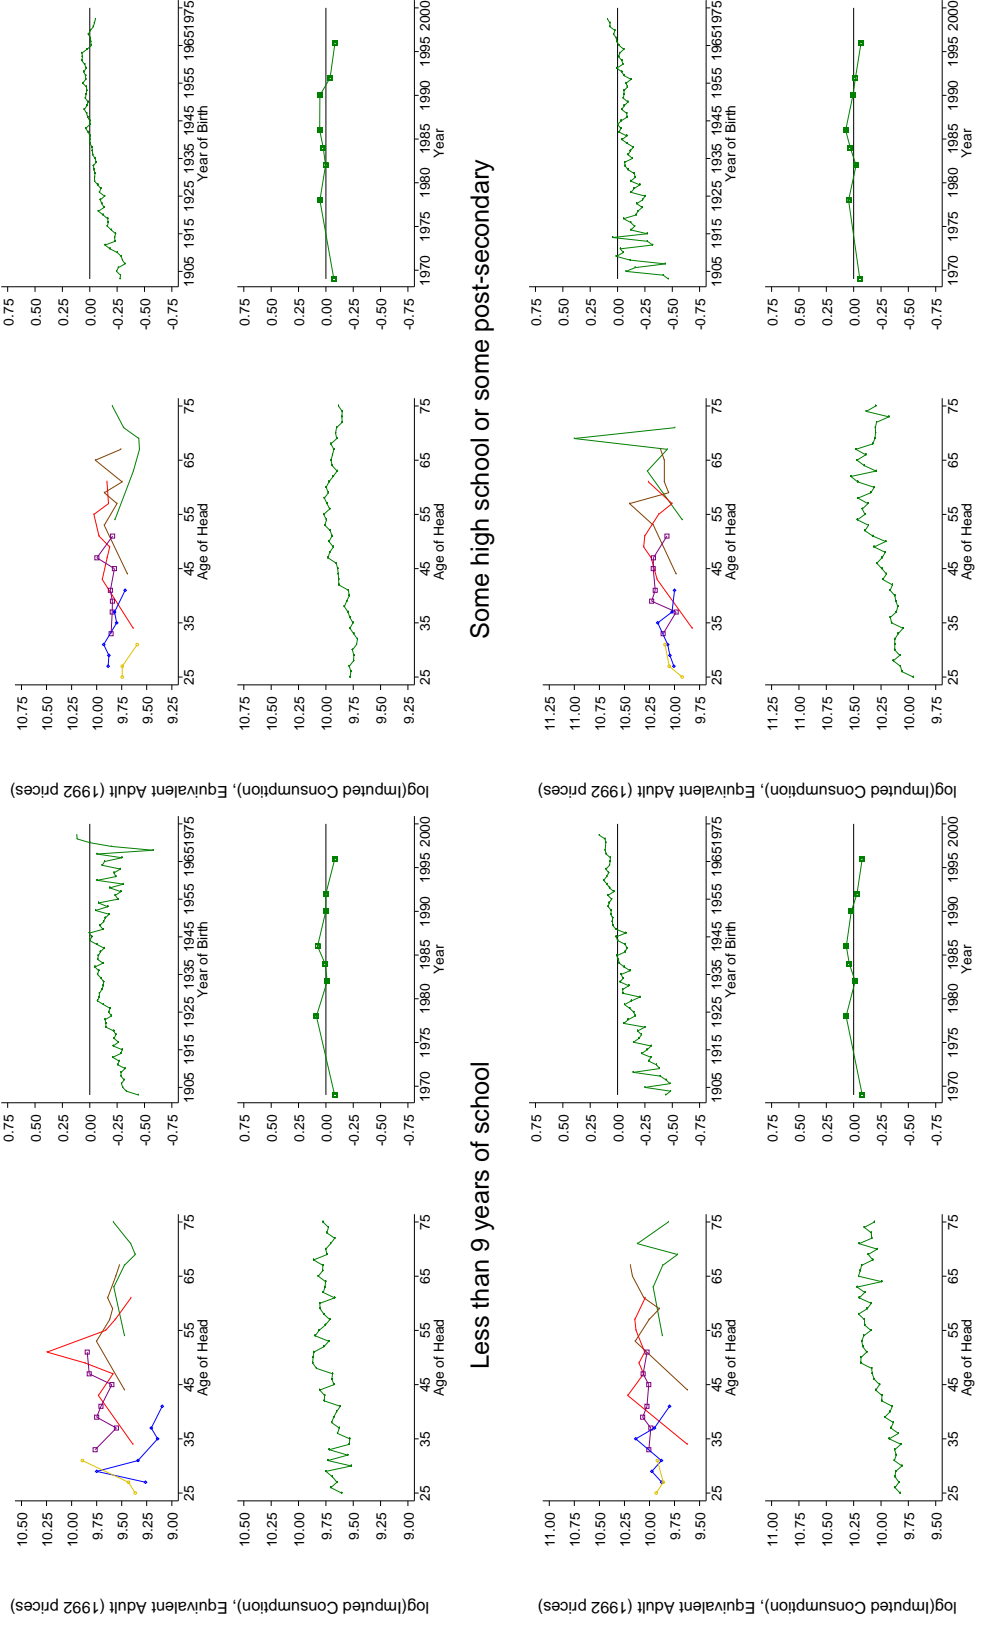


Figure 16: Imputed consumption per equivalent adult by education group




Figure 17: Household savings rate by education group




Figure 18: Share of household income from government sources by education group




Figure 19: Share of household consumption devoted to gifts by education group

Table 1: Cell Size (The cells corresponding to the cohorts born in 1944 and 1924 are emphasized in bold.)

Age	Year										All
	1969	1978	1982	1984	1986	1990	1992	1996	1997	1998	
25	275	205	185	76	127	53	94	72	107	.	1,194
26	351	211	171	112	127	72	131	109	133	114	1,531
27	312	205	188	114	140	112	134	110	141	107	1,563
28	344	201	196	104	151	117	142	108	140	130	1,633
29	308	201	198	111	171	114	149	111	143	108	1,614
30	325	245	210	94	179	131	159	133	144	121	1,741
31	310	235	179	124	163	128	155	136	180	130	1,740
32	317	243	206	112	148	132	159	152	176	141	1,786
33	292	212	176	97	164	129	166	166	199	184	1,785
34	320	187	186	110	160	107	154	188	224	179	1,815
35	325	200	205	97	146	123	186	176	215	173	1,846
36	310	165	233	95	145	116	194	192	225	185	1,860
37	320	170	153	79	168	112	147	171	241	193	1,754
38	328	190	132	102	156	99	173	207	188	173	1,748
39	335	154	133	78	150	112	187	195	201	212	1,757
40	289	144	136	88	135	102	176	156	213	211	1,650
41	319	140	120	94	116	80	171	150	223	196	1,609
42	342	148	138	68	136	109	171	199	215	154	1,680
43	310	133	97	55	118	100	139	192	190	168	1,502
44	286	145	79	59	112	96	133	190	203	175	1,478
45	300	128	77	47	113	88	140	157	188	150	1,388
46	286	139	59	54	102	82	158	142	175	160	1,357
47	277	120	90	41	94	74	130	149	182	156	1,313
48	264	132	82	42	74	61	113	147	197	142	1,254
49	327	133	81	47	67	70	104	160	167	161	1,317
50	314	152	75	35	87	71	130	129	172	161	1,326
51	227	132	71	44	97	68	90	117	128	132	1,106
52	274	133	96	51	92	87	115	125	145	128	1,246
53	293	133	81	41	66	56	87	103	143	132	1,135
54	259	143	93	47	103	48	80	103	126	104	1,106
55	277	123	78	40	87	58	102	91	126	117	1,099
56	252	124	73	43	96	70	99	87	124	108	1,076
57	247	131	71	56	95	47	77	88	115	88	1,015
58	260	149	83	43	79	54	94	84	99	91	1,036
59	220	141	78	39	96	68	85	94	107	73	1,001
60	237	127	85	39	74	62	107	72	103	79	985
61	193	112	85	50	70	52	89	86	90	82	909
62	252	119	98	48	98	55	84	76	94	94	1,018
63	224	98	87	38	67	39	86	78	106	87	910
64	205	139	85	55	67	52	103	89	87	72	954
65	263	123	86	56	90	65	78	77	102	71	1,011
66	193	124	76	34	86	47	69	104	105	72	910
67	.	114	82	54	79	57	86	92	110	73	747
68	.	106	72	40	77	52	90	85	76	74	672
69	.	115	78	44	67	53	67	68	95	66	653
70	.	99	79	59	72	57	93	.	77	73	609
71	.	84	58	42	58	52	89	.	102	92	577
72	.	81	92	39	68	42	86	.	92	75	575
73	.	88	53	31	59	40	51	.	77	75	474
74	.	75	81	37	51	44	53	.	67	66	474
75	.	81	84	43	48	43	65	.	86	82	532
All	11,962	7,432	5,790	3,248	5,391	3,958	6,020	5,716	7,364	6,190	63,071