NBER WORKING PAPER SERIES

# RETIREMENT INCOME SECURITY AND WELL-BEING IN CANADA

Michael Baker
Jonathan Gruber
Kevin S. Milligan

Working Paper 14667
http://www.nber.org/papers/w14667

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2009

This research was supported by the National Institute on Aging, Grant P01-AG005842. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Retirement Income Security and Well-Being in Canada
Michael Baker, Jonathan Gruber, and Kevin S. Milligan
NBER Working Paper No. 14667
January 2009
JEL No. H55,J14,J26

## ABSTRACT

A large international literature has documented the labor market distortions associated with social security benefits for near-retirees. In this paper, we investigate the 'other side' of social security programs, seeking to document improvements in wellbeing arising from the provision of public pensions. To the extent households adjust their savings and employment behavior to account for enhanced retirement benefits, the positive impact of the benefits may be crowded out. We proceed by using the large variation across birth cohorts in income security entitlements in Canada that arise from reforms to the programs over the past 35 years. This variation allows us to explore the effects of benefits on elderly well-being while controlling for other factors that affect well-being over time and by age. We examine measures of income, consumption, poverty, and happiness. For income, we find large increases in income corresponding to retirement benefit increases, suggesting little crowd out. Consumption also shows increases, although smaller in magnitude than for income. We find larger retirement benefits diminish income poverty rates, but have no discernable impact on consumption poverty measures. This could indicate smoothing of consumption through savings or other mechanisms. Finally, our limited happiness measures show no definitive effect.

Michael Baker
Department of Economics
University of Toronto
150 St. George Street
Toronto, Ontario M5S 3G7 CANADA
and NBER
baker@chass.utoronto.ca

Jonathan Gruber
MIT Department of Economics
E52-355
50 Memorial Drive
Cambridge, MA 02142-1347
and NBER
gruberj@mit.edu

Kevin S. Milligan
Department of Economics
University of British Columbia
#997-1873 East Mall
Vancouver, B.C.
CANADA V6T1Z1
and NBER
kevin.milligan@ubc.ca

Expenditures on income security programs for seniors in Canada are projected to increase substantially over the next decades. For example, expenditures on the primary transfer programs for seniors, which totaled over $23 billion in 1999/2000, are projected to rise from $25 billion in 2001 to $109 billion by 2030, or from 2.3 percent of GDP to 3.2 percent (Office of the Superintendent of Financial Institutions 2002a). Benefit expenditures in the largest of the two public pension plans, which totaled over $20 billion in 2000/01 are projected to rise to $74 billion in 2025 (Office of the Superintendent of Financial Institutions 2002b).

While these expenditures are growing, Canada is in a better position than other developed countries. The main earnings related pensions have moved in the direction of pre-funding and current contribution rates are projected to be actuarially stable in the future.[1] The pure pay-as-you-go component of the system funded out of general tax revenue is relatively small. Finally, population growth, while diminished relative to earlier decades, is still projected to be larger than in Europe or Japan.[2] Still, as the country's population ages, stresses on the public finances may extend to seniors' pension benefits.

There are a variety of solutions possible to this long-term problem. Some of them involve reducing, in one way or another, the benefits available to retirees in Canada. For example, the government could cut the OAS amount, change the translation of past earnings into CPP/QPP benefits, or raise program entitlement ages. As discussed in Baker, Gruber and Milligan (2003, 2004), these changes could have significant impacts on both retirement behavior and program finances. For example, we find that raising the age of eligibility for retirement

[1] See Office of the Superintendent of Financial Institutions (2002b).
[2] United Nations (2005).

programs in Canada would lower program expenditures by over 22%.

While improving program finances, such a change might also have pernicious effects on the well-being of the elderly in Canada. The income from retirement programs made up 45.5% of the average elderly family's gross income in Canada in 2000. If benefits from these programs are reduced, it could lead to significant reductions in the standard of living of elderly families.

But would it? This depends on the reaction of other sources of support to the elderly. For example, a large literature has investigated the question of whether social security benefits simply serve to crowd out savings by the elderly. Other sources of income support for the elderly include their own labor supply or transfers from other family members. If these other sources of support are increased as Social Security benefits are reduced, then there may be little implications of benefits reductions for elderly well-being.

In this paper, we investigate directly the "other side" of Income Security reform: the implications for elderly well-being. We do so by using the large variation across birth cohorts in income security entitlements in Canada that arise from reforms to the program over the past 35 years. This variation allows us to explore the effects of benefits on elderly well-being while controlling for other factors that affect well-being over time and by age.

The paper proceeds as follows. In Section I, we provide a brief background on the institutional features of the income security programs for Canadian elders. In section II, we describe our data sources. Section III discusses our empirical strategy for using variation in benefits across birth cohorts to identify the effect of benefits on well-being. Section IV presents time series evidence on the role of Income Security, and Section V presents regression evidence. Section VI concludes with the implications of our findings for Income Security policy.

# Section I: Background on the Canadian Income Security System

Public pension plans for seniors in Canada fall into two groups: 1) the Old Age Security program, and 2) the Canada and Quebec Pension Plans. The Old Age Security program, which encompasses the Old Age Security (OAS) pension, the Guaranteed Income Supplement (GIS), and the Allowance, are transfer programs financed out of general tax revenues. The Canada and Quebec Pension Plans (CPP/QPP) are contributory programs that pay a benefit which is proportional to an individual's earnings over his/her working life.

## *The Current Parameters*

OAS is the older of the two programs, dating back to 1952. Its current rules, which have been relatively stable since the 1970s, provide for a pension payable to all individuals aged 65 or older who satisfy a residency requirement.[3] In December 2005 the monthly benefit paid to individuals who fully satisfied the residency requirement was $479.83. This benefit is clawed back from higher income pensioners at a 15 percent rate, starting at incomes of $60,806 (2005). Benefits are full indexed to the CPI and fully taxable under the Income Tax Act.

Available since 1967, the GIS component of OAS is an income tested supplement that is also payable to those aged 65 or older. The income test is applied annually based on taxable

[3] Individuals must have been a Canadian citizen or legal resident of Canada at some point before application, and have resided in Canada for at least 10 years (if currently in Canada) or 20 years (if currently outside Canada). The benefit is prorated for pensioners with less than 40 years of Canadian residence, unless they are "grandfathered" under rules that apply to the persons who were over age 25 and had established attachment to Canada prior to July 1977.

income (as defined by the Income Tax Act, excluding any OAS pension) at the family level. Benefits are taxed back at a 50 percent rate, expect in families where only one partner is 65 or older, in which case the tax rate is 25 percent. There are separate monthly benefits for married and single individuals, fully indexed to the CPI, which amounted to $371.46 and $570.27, respectively, in December 2005. These benefits are not subject to income taxes.

Finally, the Allowance (formerly the Spouse's Allowance) is an income tested benefit payable to 60-64 year old partners of individuals who are aged 65 or older. It is also available to widow/widowers who are aged 60-64. This part of OAS was introduced in 1975. For the partners of 65+ year olds the benefit equals the sum of an OAS pension plus a GIS at the married rate. Benefits are reduced by 75 cents for each dollar of income to the OAS part of the benefit is eliminated. At this point the benefit (as well as the partner's GIS) is reduced according to standard GIS rules. Widow/widowers receive a slightly higher benefit the benefit, but the income test is applied equivalently. Again benefits are fully indexed and not subject to income taxes.[4]

CPP and QPP pensions, available since 1966, are financed by employer and employee contributions. In 2005, these are (each) payable at a rate of 4.95 percent on earnings between $3,500 and $ 41,100.00 (the Year's Maximum Pensionable Earnings, YMPE). Pensions are available to individuals aged 60 or older who have made contributions in at least one calendar year in the ("contributory") period stretching from age 18 to the year of application (or age 70).[5] Benefits are actuarially adjusted for applications at different ages in the interval 60 through 70.

[4] More details on the Allowance are available in Baker (2002).
[5] For individuals who were 18 prior to 1966, the contributory period starts on January 1, 1966.

The adjustment is a reduction (premium) of 0.5 percent for each month the application precedes (follows) the 65th birthday. Also, applications prior to age 65 are subject to a retirement test, which is that the individual's annual rate of earnings cannot exceed the maximum retirement pension payable at age 65, for the year in which the pension is claimed.[6]

The amount of benefits received is based on the individual's "average" earnings over the contributory period. Any months (a) receiving a disability pension, (b) spent rearing small children,[7] (c) between age 65 and the commencement of the pension[8], and (d) 15 percent of the remaining months, are excluded from the calculation.[9] In each of the remaining months, earnings are expressed as a fraction of 1/12 of the current YMPE, to a maximum of one.[10] The average of these fractions is then multiplied by the average value of the YMPE over the previous five years (including the year of application), and then by 0.25, to arrive at the benefit level. Benefits are fully indexed and subject to income taxes.

Both the CPP and QPP pay survivor pensions when a contributor dies and has made contributions for the lesser of 10 years or one third of the number of years in the contributory period.[11] For survivors under age 65, benefits are equal to a flat rate benefit plus 37.5 percent of the earnings-related pension of the deceased spouse. Benefits are reduced if the claimant is younger than age 45 and is not disabled and has no dependents. For survivors aged 65 and

[6] There are no restrictions on returning to work after the benefit is being paid.
[7] This is defined as months where there was a child less than 7 years of age and the worker had zero or below average annual earnings.
[8] Periods after age 65 to age 70 can be substituted for periods prior to age 65 if this will increase their future retirement pension.
[9] The last three of these exclusions cannot be used, however, to reduce the contributory period below 120 months after taking into account the offset for months of disability pension receipt.
[10] Excess earnings in one month above 1/12 of the YMPE may be applied to months in the same calendar year in which earnings are below 1/12 of the YMPE.

above, the pension is equal to the greater of a) 37.5 percent of the deceased's retirement pension plus 100 percent of the survivor's own retirement pension, or b) 60 percent of the retirement pension of the deceased plus 60 percent of the survivor's own retirement pension. There is an upper cap on total payments equal to the maximum retirement pension payable in that year.[12] There are also orphan benefits payable to the children of the deceased.[13]

The CPP and QPP also pay disability pensions. The benefit is equal to a flat-rate portion plus an earnings-related portion equal to 75 percent of the applicable CPP/QPP retirement pension, calculated with the contributory period ending at the date of disability. More information can be found in Gruber (2000).

## *Program Changes over Time*

Changes to the OAS and CPP/QPP programs over time are the basis of identifying any impact of income security benefits on well being. These changes have altered both the level of benefits available to claimants, as well as the eligibility criteria.

The age of eligibility for OAS pensions was changed over the period 1965 to 1970. Initially these pensions were available starting at age 70 but starting in 1965 the age was lowered annually in one year increments until it reached 65 in 1970. Full indexation of benefits did not start until 1973. Previously, adjustments to benefits were ad hoc. Finally, the residency requirements for benefits were changed in 1977. Originally individuals who did not fully satisfy the requirement received no benefit, but starting in 1977 pro-rated benefits were granted for

[11] There is also a lump sum death benefit, which is generally equal to one-half of the annual CPP/QPP pension amount up to a maximum $2,500.

partial fulfillment.

The GIS was introduced in 1967. Benefits were indexed starting in 1973. There were also one time real increases in benefits of 69 percent in 1971, 20 in 1972, 44 percent in 1979/80 and 16 percent in 1984. The introduction of the program, as well as these benefit changes, should have heterogeneous impacts across birth cohorts. For example, older cohorts who missed out on CPP/QPP pensions (introduced in 1966) are more reliant on GIS than younger cohorts whose GIS is partially taxed back because of these pensions. Also, the introduction of the Allowance in 1975 led to a dramatic temporal change in the public support available to 60-64 year olds. This support was extended to widows aged 60-64 in 1985.

The introduction of the CPP/QPP in 1966 had a differential impact across birth cohorts. The parameters of the system were phased in over a ten year "transition" period. First, pensions were initially (as of January 1967) payable starting at age 68. The minimum age was lowered in subsequent years until it settled at age 65 in 1970. Second, pensions were prorated. The benefit was calculated as the standard pension entitlement multiplied by the faction of the transition period that the individual had made contributions to the plans. For example, if the benefit application was made in January 1972 after 6 years of contributions, the standard pension entitlement would be multiplied by 0.6, which is just the number of months of contributions (72) divided by 120 months, the length of the transition period. This means that those born before 1900 got no pensions, those born between 1901 and 1910 got partial benefits, and those born in/after 1911 were 65 in 1976 and so they got full benefits. The details of the phase-in were

[12] If the surviving spouse is receiving his or her own CPP disability pension, the sum of the earnings-related portion of the two pensions cannot exceed the maximum retirement pension available in the year.
[13] More details on survivor pensions are available in Baker et al. (2004).

announced in 1965, so the 1901 cohort got a windfall, while later cohorts could have anticipated the benefits they would receive.

Starting in 1970, CPP/QPP benefits were available starting at age 65. This changed in 1984 for the QPP and 1987 for the CPP when more flexible rules were introduced. As described above, these rules allowed application anytime between ages 60 and 70 subject to an actuarial adjustment.

The real value of pensions payable by the CPP/QPP changed quite dramatically over the period 1975 and 1986. The YMPE, which is a prime determinant of benefits, was initially set to equal average earnings. However, it was not indexed, so by the mid 1970s the value of the YMPE had fallen to around 70% of average earnings. To rectify this situation, both the CPP and QPP started to increase the YMPE by 12.5% per year, stopping in 1986 when the YMPE again reached average earnings. This accounts for the large increases in the CPP/QPP benefits over the decade from the mid 1970s to the mid 1980s.

Finally, both the CPP and QPP initially had earnings tests on the pensions of beneficiaries aged 65 to 69. Benefits were recovered at a rate of 50 cents for every dollar of earnings in excess of 18 percent of the YMPE on an annual basis. An additional 50 cents was recovered for earnings in excess of 30 percent of the YMPE. Therefore, at higher levels the tax back was 100%. The earnings test was eliminated from the CPP in 1975. The second tier of the test was eliminated from the QPP in 1973, while the remaining tier was dropped in 1977.

In Figure 1 we graph the real value of the different pensions and benefits starting in 1961. Many of the reforms and changes reviewed can be clearly seen in this picture. In particular, the real increases in the GIS in 1971/72, 1978/79 and 1984, the indexing of the OAS pension, the

effects of the CPP transition period and the recovery of the YMPE to the level of average wages from 1975 to 1986.

## Section II: Data Sources

Our analysis encompasses three aspects of elderly wellbeing: incomes, consumption, and happiness. In all cases, we use survey microdata provided by Statistics Canada. Below, we describe the details for each of our data sources in detail. In addition, we explain exactly how we form the wellbeing variables we use in the analysis from the raw data sources. Table 1 provides descriptions of our data sources and the variables we form.

There are several issues common to all of the data sets that we use. First, we use 2002 Euros for all of the analysis. We update dollar values to 2002 using the Canadian Consumer Price Index, then translate to Euros using the average 2002 Euro-Canadian dollar exchange rate (1.4832). Second, we scale all income and consumption data by an equivalence scale to account for the size of the household / family. The first adult is counted as one, each subsequent adult counts as 0.7, and each child under age 18 counts as 0.5. A third issue is the definition and weighting of elderly households. We define an elderly household as one in which there is at least one member age 60 or over. All other households are therefore classified as non-elderly. Furthermore, we weight the results for the elderly households by how many members are ages 60 or over, effectively placing our results on an individual basis for the elderly.

## *Income data*

We draw our income data from two different sources. From 1971 to 1997, we have data from the Survey of Consumer Finances. This survey reports information on the education, age, family structure, labour market activity, and incomes of respondents. In particular for income, the data is disaggregated to a degree that serves our purposes well. In particular, we use the variables for after-tax income, Canada / Quebec Pension Plan income, and OAS/GIS/SPA income to measure the actual public pension benefits received by elderly Canadians.

The Survey of Consumer Finances is conducted using the sampling frame of the monthly Labour Force Survey, which is a stratified random sample of Canadians with some oversampling of smaller provinces to ensure sufficient sample size. With the survey weights, nationally representative results can be obtained.[14] The survey is available in odd years from 1971 to 1981, then 1982, and finally on an annual basis from 1984 to 1997. The survey is available for different family definitions and for individuals. We use the census family sample, which ranges in size from 25,927 up to 47,840.[15] Unfortunately, the household version of the survey does not contain disaggregated government pension income, so we must use the family version.

The Survey of Consumer Finances was replaced in the late 1990s by the Survey of Labour and Income Dynamics. The new survey is available for 1993 to 2002 and for our purposes is similar to the previous survey. One interesting difference is the source for income

[14] The sampling frame excluded residents of the three northern territories, those living in institutions, and inhabitants of native reserves. These groups typically account for less than three percent of the population.
[15] A census family is comprised of parent(s) with their children, childless couples, or grandparents living with their grandchildren. Individuals on their own who are not in a census family are classified as "non-family individuals", but are included in the census family sample.

data. For the Survey of Consumer Finances, income data was asked on a recall basis. For the Survey of Labour and Income Dynamics, however, respondents may check a box indicating consent to attach data directly from tax records. A high percentage of respondents take this option (over 80 percent in 2000). We append the 1998 to 2002 data from the Survey of Labour and Income Dynamics to the Survey of Consumer finance data for 1971 to 1997 to complete our income data.

We form several income variables for our analysis. The first is total income from public pensions, which is the sum of reported CPP/QPP and OAS-GIS-SPA income. Second, we take the reported total after-tax income for the family. We use the mean, along with several percentiles of this variable. We also use the mean calculated just for those who are beneath the poverty line. Third, we form a measure of relative income poverty. For each year, we find the median family income for non-elderly families. We set a poverty line at 40 percent of this median, then form a dummy variable for each elderly family indicating whether it is over or under the poverty line. Fourth, we measure absolute poverty by repeating the same exercise but using the poverty line for 1971 in all years. Combined, these four measures will richly describe the income patterns of elderly Canadian families over the three decades we study.

### *Consumption data*

The consumption data we use come from two different surveys, each conducted over a number of years. The first survey is the Family Expenditure Survey. It was first conducted in 1969, followed irregularly by surveys in 1974, 1978, 1982, 1984, 1986, 1990, 1992, and 1996. The sampling frame in several of the years included only residents of large cities, so we restrict our sample in all years to residents of large cities to maintain comparability. This survey is also

based on the Labour Force Survey sampling frame, so survey weights can potentially produce nationally representative results. The sample size varies considerably, with between 4,569 and 15,140 observations.

The unit of observation changed for the Family Expenditure survey in 1990. Prior to 1990, the concept used was the spending unit, defined as "a group of people living in the same dwelling who depend on a common or pooled income for major expenses or one financially independent individual living alone." This is similar to a definition of an economic family, one or more of which can live in any one dwelling. From 1990 on, the unit of observation was the household, defined simply as those in the same dwelling unit.

The second set of consumption data is the annual Survey of Household Spending, available for the years 1997 through 2002. This survey has larger sample sizes than the Family Expenditure Surveys did, with around 14 and 18 thousand households per year. We discard the 2002 survey because single years of age are not reported in the public-use version of the data.

The key variable we extract from these surveys is the current consumption of the household. Included in this measure is spending on goods and services throughout the year. It is equivalent to the total expenditure of the household less personal taxes, gifts and donations, and life insurance. We use the mean, as well as several percentiles of this variable. In addition to the current consumption variable, we construct two other measures of elderly wellbeing similar to those we made for income. There is a relative consumption poverty measure constructed as being under 40 percent of the median non-elderly household's consumption level in each year, and an absolute measure based on the median non-elderly household in 1969.

### *Happiness data*

To undertake our analysis of the happiness of elderly Canadian households, we put together several waves of the General Social Survey. The sample is formed using a random digit dialing methodology, with supplemental samples drawn from the Labour Force Survey sampling frame in some years. Sample sizes are around ten thousand observations per year. The happiness question that we require is asked in the years 1985, 1986, 1989-1991, 1996, and 1998. Unfortunately, age is only reported in 5-year ranges, so we cannot identify the exact year of age of the respondents. As the happiness question is individually based, we use the individual as the unit of observation for the happiness data.

The exact wording of the question is "presently, would you describe yourself as . . ." Possible responses are very happy, somewhat happy, somewhat unhappy, and very unhappy. We form two variables to use for our analysis. First is an indicator for having responded "very happy." The second is an indicator for having responded either somewhat or very unhappy.

## Section III: Empirical Strategy

For the regression analysis that forms the core of this paper, we want to regress well-being outcomes on retirement income. Because observed retirement income may be determined by the same factors that determine the outcomes we wish to study, we create *simulated benefits* that are exogenous to the outcomes. We do this by attempting to capture legislative variation in the system over time. In general, benefits are formed from two components: laws and the

characteristics of recipients. The laws come in two forms: the euro amount of the benefits and the age at which benefits are available. The idea of the simulated benefit approach is to abstract from all differences in characteristics of recipients and focus solely on variation in benefits that arises from law changes. Ideally, we would take exactly the same person, put him in every birth cohort, and then compute his benefits. In this way, we would hold characteristics constant, and any benefits variation that we saw over time or across birth cohorts would be due to law changes.

Of course, this is impossible in reality. Two types of factors may differ across birth cohorts that can affect benefit determination. The first type includes factors which are unlikely to be endogenous to Income Security rules (at least to a first approximation), but which are potentially important omitted determinants of well-being: differences in earnings histories; differences in capital income (which can matter for GIS eligibility and benefits); and differences in spousal labor supply. The second type is factors which are likely endogenous to Income Security rules, such as differences in actual observed retirement ages. One possible simulated benefits approach would hold the first type constant, and compute benefits for each birth cohort based on that constant comparison. A more difficult question is whether one wants to hold the second set of factors constant, as this may be part of the effect that we want to capture in our results.

Thus, we consider three approaches. First, we show results using actual benefits received, as measured by reported IS income at each age. Second, we employ a "partial simulation" approach, in which the first set of factors (earnings, capital income, family status) is held constant, but the second set of factors (retirement ages) is allowed to vary. In particular, we base the benefit amount on a fixed earnings history across all birth cohorts, not the actual earnings

history by cohort. The motivation for this partial simulation approach is to take account of changes in retirement behaviour that may be responses to changes in program structure. At the same time, by using the fixed earnings history the mixed approach allows for the fact that there could be correlated changes over time in tastes for leisure (as manifested in early retirement) and measures of well-being that could bias the results if each cohort's actual earnings were used. Finally, we also consider a "full simulation" approach, whereby both the first and second sets of factors are held constant. For that approach, we use the base cohort earnings history *and* the base cohort retirement patterns. So, in the full simulation, the variation comes only from legislative variation since earnings and retirement variation are held constant for the entire sample.

Our empirical approach exploits the available policy variation. We concentrate on age-year cells. By calculating the expected benefits for a typical person in each age-year combination, we can exploit changes that occur through time (like the GIS and CPP changes) that affect people of different age groups differently. For example, the GIS expansions affected those ages 65 and older, but not those who were younger. The focus on age-year cells, therefore, provides both a clear and an effective source of identification.

## *Methodology*

In order to impute benefits to a retired family we must know their earnings history, their capital and private pension income, at what age they retired, their marital status, and the rules that were in place that determined their benefits. Combined, this information allows us to calculate 'typical' benefits for families in our data. Our empirical strategy involves holding

several factors constant across all of the cohorts in our sample. We do this by focusing on the 1920 birth cohort, which reaches age 65 in the middle year of our sample (1985). By pushing this cohort through the benefit calculation for all years in the sample, we will be able to extract only the variation derived from changes in policy rules, since the cohort characteristics will be identical by construction.

For the age-year analysis, we form the earnings profile by taking deciles of earnings in each year and age in the Survey of Consumer Finance data over all individuals, separately for men and women. We perform this calculation only for the 1920 birth cohort, which we will use to simulate benefits for all cohorts. We then average these deciles to form an earnings value for each age and year. The motivation for this decile-based approach is to capture non-work in the best way possible. Median earnings of women in earlier years were zero, for example, so using the decile approach allows us to capture the fact that at least some of the women in those years had earnings. Since the oldest income microdata available to us is 1971, we impute earnings to each year previous to 1971 by deflating the 1971 age-earnings profile by the growth in the industrial composite wage. For calculating CPP/QPP benefits, we need to extend back to 1966.

For capital income, we use data from the Survey of Consumer Finances. We define capital income as the sum of income from investments and income from retirement pensions, which includes both withdrawals from Registered Retirement Savings Plans and Registered Retirement Income Funds. For the age-year analysis, we form deciles of capital income for each age and year by sex. The decile approach allows us to capture nonlinear effects of the income security system that would be missed if we simply used the median or mean for assigning capital income. Again, we use only the 1920 cohort's capital income profile, adjusted up and down for

other cohorts using the CPI. We assign the capital income deciles to the earnings deciles by assuming that the lowest earning decile has the lowest capital income, the 2nd lowest earners with the 2nd lowest capital income, and so on. At the end of the calculation, we average over the deciles to arrive at the simulated benefit.

We do not observe the age of retirement for elderly Canadians in our data. Because CPP/QPP benefits depend on the date of retirement through the actuarial adjustment of early retirement benefits, we must account for the possibility of different retirement dates. We do so by calculating the retirement benefits that would be received at each potential age of retirement between the ages 55 and 70. We then average these benefits using the observed retirement probabilities at each age for the cohort as weights. We form the cohort retirement probabilities with the Survey of Consumer Finances data. We assume that none are retired at age 55 and that all are retired by age 70. The retirement rate is the change in the proportion of respondents who are employed in one year over the previous year. Using these observed rates, we form a cohort- and sex-specific set of probabilities for retirement that sum to one. With these probabilities, we then determine a retirement age-weighted average of the benefits the family is currently receiving.

Finally, we must consider different family types. Someone we observe at an older age may have previously been married but is now widowed. For such a person, we must assign spousal benefits in order to get a correct measure of total Income Security benefits. We take observed family types by age and sex in the 1990 General Social Survey and develop a set of

probabilities.[16] For each family in our simulation, we average the benefits over each of the seven family types using the 1990 probabilities as weights.

Our goal is to have simulated benefits for each someone in each age-year cell. To form the full and the partially simulated benefits we use for the analysis, we use the described methodology in slightly different ways. For both simulations, we use the 1920 cohort earnings profile, the 1990 cross-sectional family types, and the 1920 cohort capital income profile. For the partially simulated benefit, we use each cohort's observed retirement pattern to find the simulated benefits. In contrast, the fully simulated benefit applies the 1920 cohort's retirement pattern to all cohorts in our data. So, the key difference between the partially and fully simulated benefits is whether cohort-specific or fixed retirement rates were used in the construction of the benefits.

## *Description of Variation*

The result of this simulation exercise is a benefits measure that varies across birth cohorts only for the age-year cells. In the regression analysis below, we will control for both age and year effects, allowing us to identify the impact of income security programs on well-being solely from variation across birth cohorts. In this section, we show that such variation is sizeable, and explain its legislative origins.

Figure 2 opens our discussion by showing the evolution of IS income over time. The first line in this figure shows the average benefits paid to elders under IS programs, in 2002 Euros per person. The line is derived from aggregate spending on the programs divided by the number of

[16] The seven family types we consider are married, male single never married, male widowed, male

Canadians age 60 and over.[17] The rise has been substantial, from around 2,000 Euros in the early 60s up to 3,000 in the 1970s and then 6,000 by the 1990s. The increase from year 2000 over 1965 is 263 percent. The second line shows the ratio of those average benefits to average worker income, with the scale on the right-hand side axis. Over the same 1965 to 2000 period, this ratio increased by 179 percent, indicating that benefit growth has far-outpaced earnings growth over the last generation. This gives an indication of the very large extent to which Canada's programs have been expanded.

Figure 3 then shows our partially and fully simulated benefits measures, graphed over time. These also grow steadily over time, in parallel fashion to actual benefits income, and move very closely together. While they follow the same trends, our simulated benefits appear slightly lower than actual benefits for the time period under consideration.

Figure 4 shows the evolution of these benefits measures over birth cohorts more clearly. We graph the simulated benefits levels for each cohort in the year they reach age 65. The first upward increment in the late 1960s corresponds to the extension of OAS benefits down to age 65. The great rise in the 1970s is caused by the phasing in of the CPP/QPP over 1970 to 1976, the introduction and expansion of the GIS in 1967, 1971, 1979, and 1980, and finally the SPA in 1976. Through the 1980s and 1990s, benefits grow much more slowly, increasing only with higher earnings through the YMPE. The partially simulated benefits line is more volatile than the fully simulated line, reflecting the additional variation introduced by the cohort-specific retirement rates.

---

separated/divorced, female single never married, female widowed, female separated/divorced.

[17] The source is Human Resources Development (1999).

The final picture of the variation in Canada's recent history is in Figure 5. Here, we expand on the age 65 variation from the previous figure to show benefits at several ages. Benefits at age 60 are zero until the 1980s when early retirement is introduced in Quebec (in 1984) and the rest of Canada (in 1987). Benefits at ages 70, 80, and 90 are similar until the mid 70s, when those who are 70 begin to show CPP/QPP income. These same cohorts a decade later as 80 year olds begin showing CPP/QPP income in the 1980s, so they pull higher than the 90 year olds who do not receive any CPP/QPP.

## *Regression Framework*

Having created these various measures of benefits (actual, partially simulated, fully simulated), we can then use them in regression analysis to assess how income security generosity affects our various measures of well-being. We make use both of the partial and the fully simulated benefits. All regressions use age-year cells as the unit of observation. We weight by the cell size from the micro-data in order to produce results that correspond to the full survey outcomes.

We begin by looking at 'first-stage' analysis of the effects of our simulated benefits on the actual benefits. In the specification, we include the simulated benefit, age dummies, year dummies, and a vector of cohort characteristics $X_{ay}$. The characteristics are measured for the cohort at age 50, and include their average earnings, education, and proportion married.

$$ActualBenefits_{ay} = \beta_0 + \beta_1 SimulatedBenefits_{ay} + \beta_2 AGE_a + \beta_3 YEAR_y + \beta_4 X_{ay} + e_{ay},$$

For some of the results we use the mean of the simulated benefit as the policy variable on the right-hand side. However, for some of the poverty measures, using a more targeted percentile of the within-cell distribution is more appropriate. We make clear what measure of simulated benefits is used when we discuss each regression result below.

We follow the 'first-stage' regressions with results that examine the 'reduced form' effect of the simulated benefits on the well-being outcomes. The specification is similar to the first-stage specification above.

$$Outcomes_{ay} = \beta_0 + \beta_1 SimulatedBenefits_{ay} + \beta_2 AGE_a + \beta_3 YEAR_y + \beta_4 X_{ay} + e_{ay},$$

The simulated benefits are either the partial or the fully simulated benefits. Again, we also use different measures of the benefits within-cell for the different outcome measures.

Our third and final specification uses a full instrumental variables approach. We use the simulated benefits to predict actual benefits, then use the predicted actual benefits as a regressor for the outcome variables. So, the simulated benefits act as an instrument for the observed actual benefits.

$$Outcomes_{ay} = \beta_0 + \beta_1 ActualBenefits_{ay} + \beta_2 AGE_a + \beta_3 YEAR_y + \beta_4 X_{ay} + e_{ay},$$

# Section IV: Evidence

## *Time Series Evidence*

We begin our presentation of the effects of IS programs on well-being by examining the evolution of our various well-being indicators over time. We use the equivalence-scale adjusted

measures, so the data represent the amount per effective person in 2002 Euros. In each case, we show two lines on each graph. One line shows the evolution of the outcome in question for the elderly in our sample, while the other provides the same measure for the working-age families in order to provide a basis for comparison. We normalize each of these lines to 100 in the first year the outcome is available, in order to emphasize the percentage change through time.

Figure 6 shows the graph for actual benefits received, as reported in our microdata sources. Similar to the aggregate and simulated benefits shows in earlier figures, there is a strong increase from 1971 through 2000, with a slight flattening out over the last few years. There is no non-elderly line in this graph since non-elderly do not receive retirement benefits.

Figure 7 shows the results for average after-tax family income. Average incomes rise by about 40 percent from 1970 through 1980, and very slowly thereafter, for both the elderly and the non-elderly. Both lines jump for the 1998 to 2002 period. Two factors underlie this break in trend. First, it reflects the change from the Survey of Consumer Finances over to the Survey of Labour and Income Dynamics. The difference is in the average reported size of census families across the two data sources – families in the new survey are smaller on average than in the older survey, leaving the income to be spread less thinly. However, there is also a sharp increase in incomes noticeable in 1998 when one compares data across the available years of the Survey of Labour and Income Dynamics. So, some of this increase does reflect growing incomes.

Figures 8 and 9 show the evolution of poverty over this time period, with the relative and absolute income poverty measures respectively. By both measures, there is a steep reduction in elderly poverty relative to the non-elderly over the 1970s and 1980s, which is consistent with IS program growth. For relative poverty, however, this reduction does not begin until the late

1970s, after the major CPP/QPP expansion has completely phased in and the GIS expansions of the late 1970s have been implemented. We can see clearly the drastic drop in elderly poverty, from 100 in 1971 down to under 40 by the early 1980s, for a drop of 60 percent. For absolute poverty, the reduction is very steep from 1971 onwards for the elderly, but it is also steep for the non-elderly as well. Once again, the major improvement for the elderly relative to the non-elderly is in the late 1970s and early 1980s. The timing of these movements corresponds well with the GIS and CPP expansions.

Figure 10 graphs the evolution of family after-tax income by percentiles. We show the $10^{th}$, $50^{th}$, and $90^{th}$ percentiles for the elderly, along with the $50^{th}$ percentile for the non-elderly. One line is sharply different than the others – the $10^{th}$ percentile of elderly income. It rises to around 230 by the 1990s, for an increase of 130% relative to 1971. In contrast, the other measures increase to only 140-160. Again, this reflects the phasing in of the CPP/QPP along with the expansions of the GIS.

We next turn to our consumption-based measures of well-being. Here the time series are less complete due to more infrequent surveys. The volatility reflects in part the differences in survey coverage and sample size across different waves of the consumption surveys. Nonetheless, it is clear from Figure 11 that the time series for the elderly and non-elderly consumption largely move together, suggesting relatively little effect of IS programs.

Figures 12 and 13 examine relative and absolute consumption poverty measures. In both cases, the elderly and non-elderly move closely together until the late 1980s and early 1990s, at which point there is a larger fall in elderly consumption poverty than non-elderly consumption poverty. Once again, given the more or less steady rise in IS benefits over this period, it is hard

to attribute this to IS effects.

Figure 14 graphs the percentiles of consumption, as we saw in Figure 10 for income. There appears to be little difference in the increase for the $10^{th}$, $50^{th}$, or $90^{th}$ percentile of consumption for the elderly, as they are all very close to 130 by 2001. The consumption of the non-elderly is slightly higher. This is in sharp contrast to the income graph in Figure 10 that showed a huge increase in income at the $10^{th}$ percentile.

Finally, Figures 15 and 16 show the results for our happiness measures. There is clearly variation across different samples, reflecting either true differences in happiness across time or some kind of differences in the sampling methodology. The measures for the elderly and non-elderly move very closely, suggesting little effect of the IS expansions on elderly happiness.

## *Regression Results*

We next move from time series graphical analysis to regression analysis, which allows us to exploit the variation in benefits across birth cohorts documented in Figure 5. As is clear from that figure, a major source of variation is the evolution of benefits at age 60 over time, relative to other ages. This allows us to include both age and year dummies in our model, controlling for general differences in well-being by age and over time. Figure 17 shows our measure of simulated benefits graphed against income, by age/year cell. It is clear that the set of points to the left, for 60-64 year olds, is described by a different process than is the set of points to the right, for 65 and over. The age 55-59 group is even more concentrated along the axis, as they receive no benefits under the simulation. Given the differences between the 60-64 and the 65+ groups, we conclude that including age dummies is the proper specification rather than trying to

fit a linear age term across both clouds of data.

The results of our analysis are presented in Table 2 and Table 3. In Table 2 for each outcome variable we show the mean and the number of observations. Across the next 6 columns are the results for the first-stage, reduced-form and instrumental variables analysis. Table 2 presents the results for eleven measures of income.

The first set of results focuses on the first stage analysis. Actual income is regressed on simulated benefits. The first result uses the means of simulated benefits to predict the mean IS income. The coefficient of 0.759 indicates that a one euro increase in partially simulated benefits leads to a 75.9 cent increase in after-tax income. For the fully simulated benefits, the coefficient is slightly higher, at 0.830. Both of these are highly statistically significant.

The next row in Table 2 shows the results for the mean income of those under the poverty line. The simulated benefits used in these regressions are the $10^{th}$ percentile of benefits. The results are negative and only marginally statistically significant, suggesting some evidence that higher benefits leads to a decrease in income. A contributing factor to the odd results is the fact that elderly poverty is very small after the CPP/QPP became fully phased in and the GIS was expanded, as was made clear in Figure 8. It is likely that most of those in poverty would be those who were not eligible for the GIS or the CPP/QPP, such as newer immigrants. Since we can't control for these factors in the regression, it might be these excluded characteristics that drive the result.

The next three rows show the results for estimates at different points of the income distribution. We have three dependent variables, formed as the mean of IS income between the $5^{th}$ and $15^{th}$ percentile, the $45^{th}$ and $55^{th}$ percentile, and the $85^{th}$ to $95^{th}$ percentile. The simulated

benefits we use in the three cases are the 10th, 50th, and 90th percentiles. The results indicate that the simulated benefits are strong predictors of income at the 10th and 50th percentiles, but do worse for the 90th percentile. Because higher income seniors likely have a larger proportion of their incomes coming from non-governmental sources, the lower predictive power of the simulated benefits at the 90th percentile should not be surprising.

The bottom half of Table 2 presents the reduced form and instrumental variables estimates for the well-being outcomes. The first outcome is mean income. Using the fully simulated benefits, the reduced form specification indicates that a one dollar increase in benefits leads to a 89 cent increase in mean income. When we use the simulated benefit as an instrument, the predicted impact of a dollar increase in benefits is 1.073. This suggests that benefits do not crowd out other sources of income. In contrast, the partially simulated benefit gives a result of 0.546, which does suggest some evidence of crowd-out.

The next two rows show the results for poverty. We use the simulated benefits of the 10th percentile of lifetime earnings for this estimation, since it is more relevant for the estimation of poverty rates. The results for partially and fully simulated benefits are quite similar. For relative income poverty, an increase in benefits of €1,000 is predicted to decrease poverty by 2.6 percentage points for the fully simulated IV specification. The absolute poverty measure shows no statistically significant effect.

To gauge the magnitude of the relative poverty coefficient, we can compare the trend in poverty over our sample period to the changes in benefits to see what proportion of the change can be explained by the benefit changes. In the age-year cell data, the 1st decile of benefits averaged €1,137 for those 60 and over in 1971, rising to €3,847 in 2002 for an increase of

€2,710. Our estimate of -0.026 implies that poverty should drop by about 7.10 percentage points. In the data, actual poverty dropped by 7.38 percentage points over this period, meaning that the increase in benefits can explain 96.3% of the trend. In other words, the decrease in poverty over the 1971-2002 period can almost entirely be explained by the increase in the generosity of benefits for low-income seniors.

The final three rows show the results at different percentiles of the income distribution. At the $10^{th}$ percentile, both the partially and fully simulated benefits reveal little evidence of income crowd out, with IV coefficients close to one. At the $50^{th}$ percentile, the partially simulated benefit coefficient is 0.330 and the fully simulated benefit is 0.728. For the $90^{th}$ percentile, the partially simulated result is not significant, and the fully simulated result comes in at 2.502. At the higher income levels, benefits represent a lower proportion of income which may make the results more sensitive.

In Table 3, we repeat the reduced form and instrumental variable analysis with our consumption and happiness measures. The first row shows the results for mean consumption. The IV results are 0.660 for partially simulated and 0.797 for fully simulated benefits. This suggests some crowd out of consumption. The next two rows look at consumption poverty. For none of the specifications – either for relative or absolute poverty – do we find a significant coefficient.

Taking these results along with the income poverty results together, an obvious question is how income poverty can be relieved by the IS system but consumption poverty shows no effect. To resolve this puzzle, we turn back to some graphical analysis in Figure 19. We plot income

and consumption poverty rates in 2000 by age. The income poverty line shows a very sharp run-up at ages 61 to 64. The average retirement age in Canada is around age 61, but the GIS and OAS benefits do not begin until age 65. The graph suggests that the lack of benefits in the 61-64 age range is leading to substantial income poverty – reaching over one fifth of the population. However, the consumption poverty line shows almost no impact over the very same age range. This may indicate that the seniors are using savings or other channels of adjustment in order to smooth their consumption while awaiting the onset of their income security benefits at age 65.

The next three rows display the results at different percentiles of the consumption distribution. The partially simulated benefits are not significant in the IV regressions for any of the $10^{th}$, $50^{th}$, or $90^{th}$ percentiles. For the fully simulated benefits, both the $10^{th}$ and $50^{th}$ percentile results are significant at the 10 percent level, showing an effect of benefits of 0.383 and 0.275 respectively. The 90t percentile result is not statistically significant.

The final two rows of Table 3 examine the effect of benefits on measures of happiness. For the very happy question, we see no sign of a statistically significant relationship between benefits and being very happy. On the other hand, there is some evidence of a decrease in reports of being unhappy or very unhappy with higher benefits in the reduced form results, but not in the IV results. These weak results may be due to the lack of variation in the data we have for happiness – the data is much more limited than the consumption and income data.

## Conclusion

We find strong evidence that Canada's income security programs have improved the welfare of the cohorts in our data. Our estimates suggest that these cohorts have more income

and consumption than they would in the absence of income security. For poverty, we find large decreases in relative income poverty, suggesting that the expansions of income security programs in the 1970s had a substantial impact on poverty. However, we do not find evidence that consumption poverty changed, suggesting that in the absence of benefits, poor families may have found some way to consume at adequate levels. Finally, we find inconclusive evidence on measures of happiness.

As a caveat, extrapolating these results beyond the observed cohorts must be done with care. Policy changes may affect savings and consumption behaviour of generations that had more time to alter their plans than for cohorts who were surprised by policy changes in their near-retirement years.

**References**


Baker, Michael, "The Retirement Behaviour of Married Couples: Evidence From the Spouse's Allowance", *Journal of Human Resources*, 37, Winter 2002, 1-34.

Baker, Michael, Jonathan Gruber, and Kevin Milligan (2003), "Simulating the Response to Reform of Canada's Income Security System," NBER Working Paper No. 9455.

Baker, Michael, Jonathan Gruber, and Kevin Milligan (2004), "Income Security Programs and Retirement in Canada," in Jonathan Gruber and David Wise (eds.) *Social Security Programs and Retirement Around the World: Micro Estimation*. Chicago: University of Chicago Press.

Human Resources Development (1999) HRDC Redbook. Statistics on income security.

Office of the Superintendent of Financial Institutions (2002a), *Actuarial Report on the Old Age Security Program as at December 31, 2000.* Ottawa.

Office of the Superintendent of Financial Institutions (2002b), *Canada Pension Plan: 19th Actuarial Report as at December 31, 2000, Supplement.* Ottawa.

United Nations (2005), "World Population Prospects: The 2004 Revision," United Nations Department of Economic and Social Affairs, Population Division. New York.

**Figure 1: Real Maximum CPP, OAS and GIS Benefits 1961-2002**

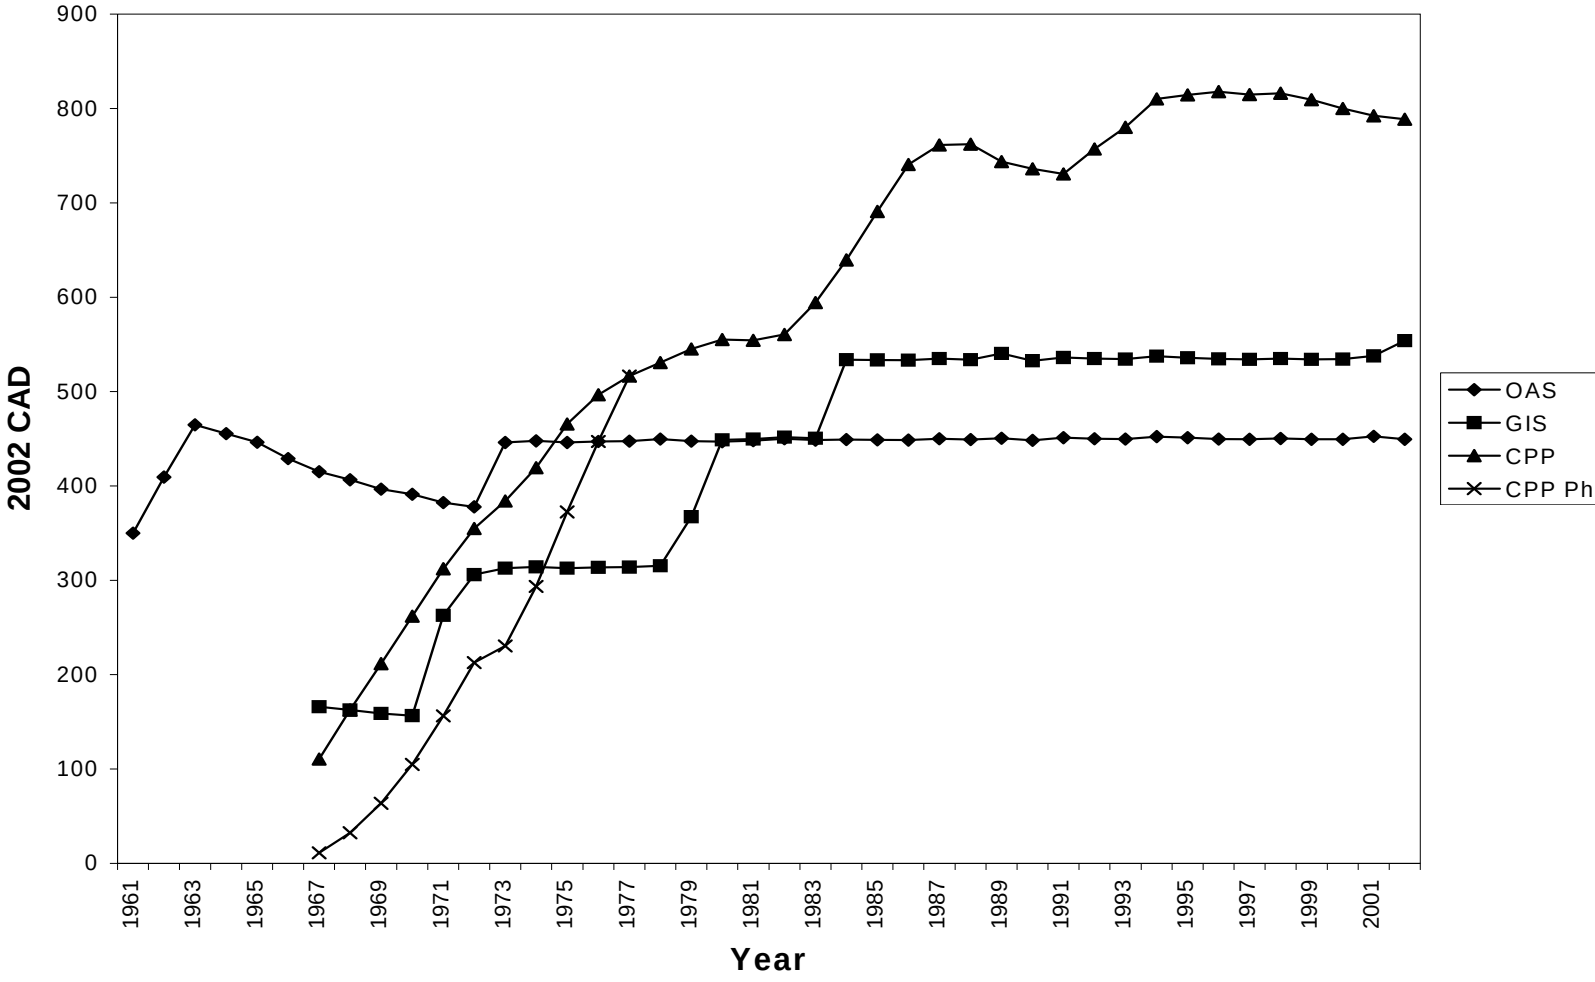


*Notes:* GIS benefits are for a single individual. The CPP phase-in calculation is for December of the indicated year.

**Figure 2: Average Income Security Benefits**

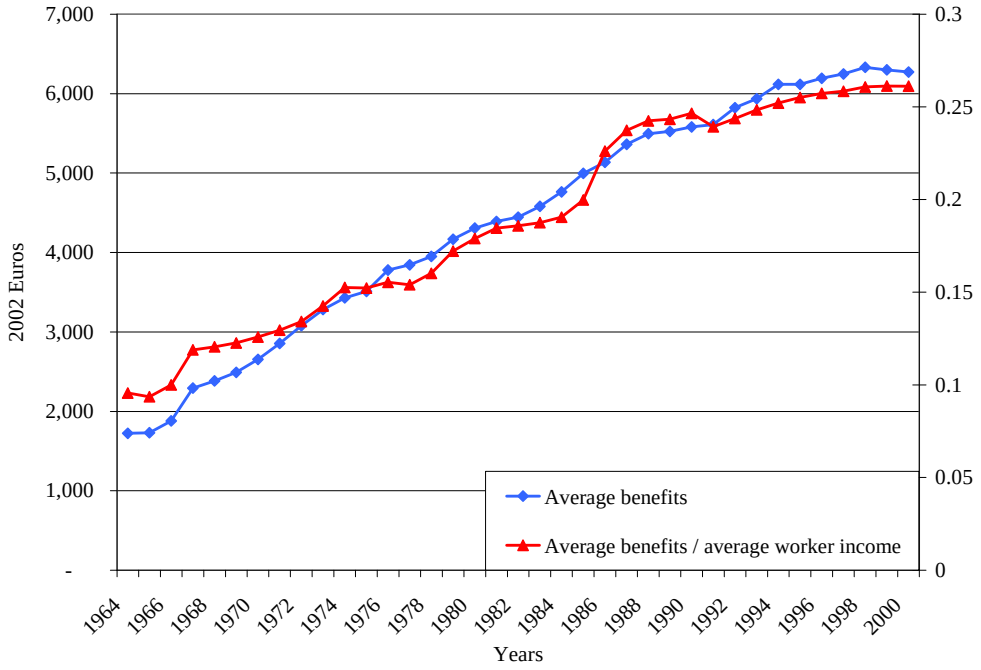


**Figure 3: Average Simulated Benefits**

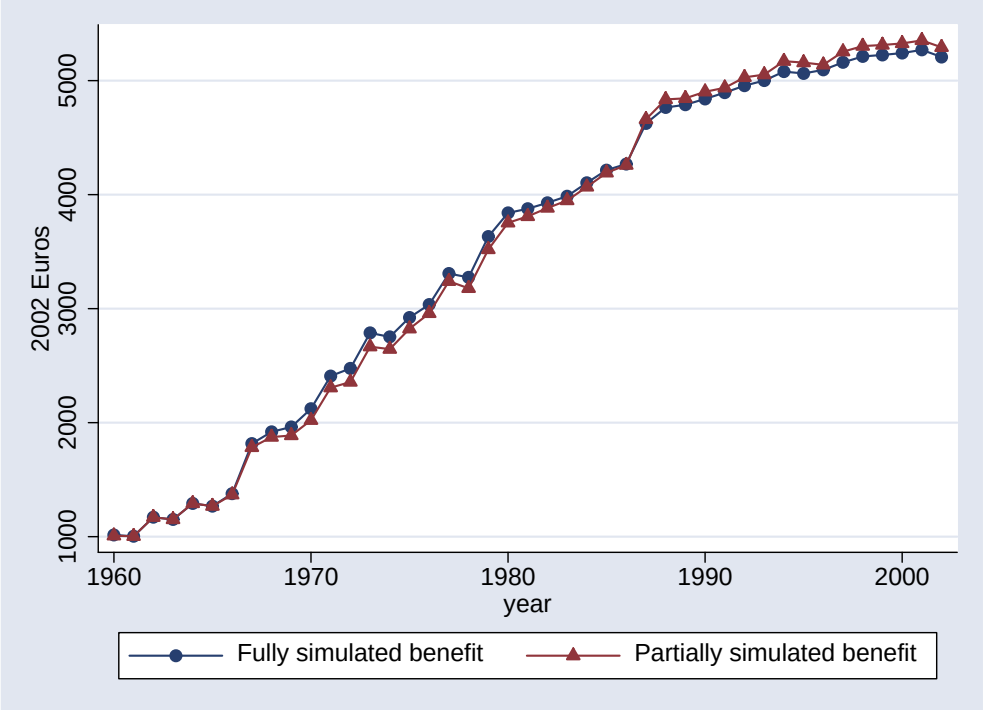

**Figure 4: Benefits at age 65 in different years**

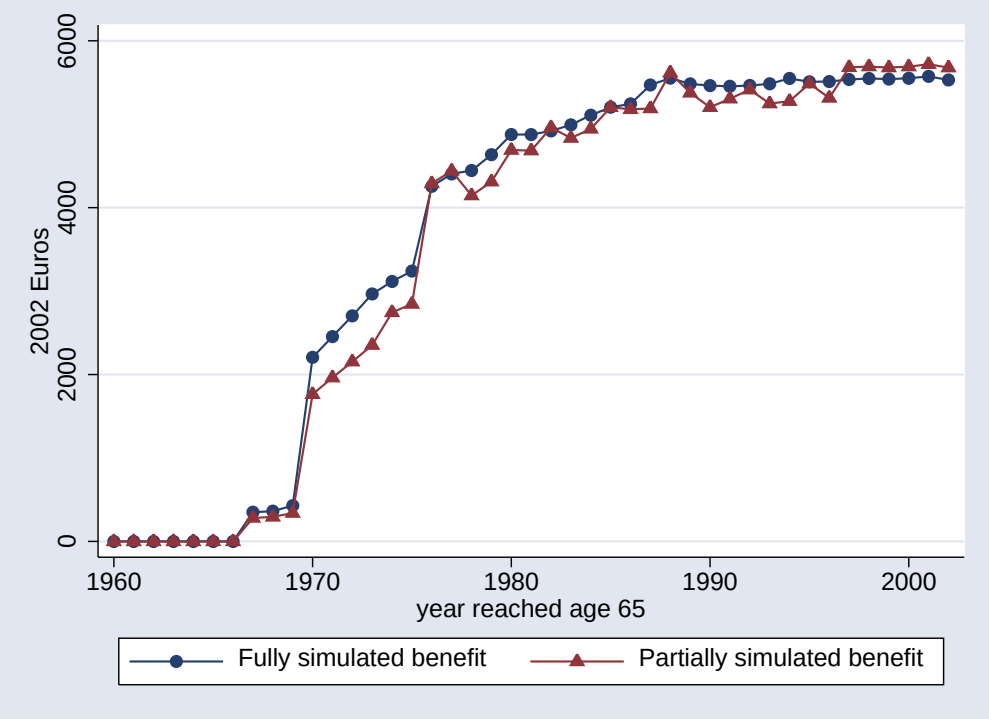


**Figure 5: Benefits at different ages through time**

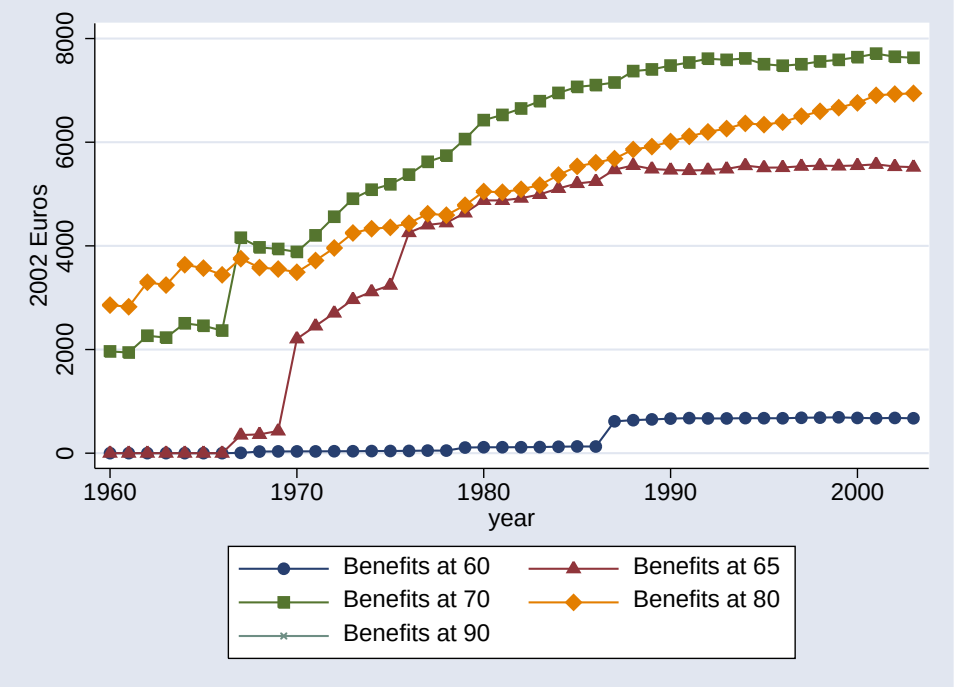

**Figure 6: Benefits**

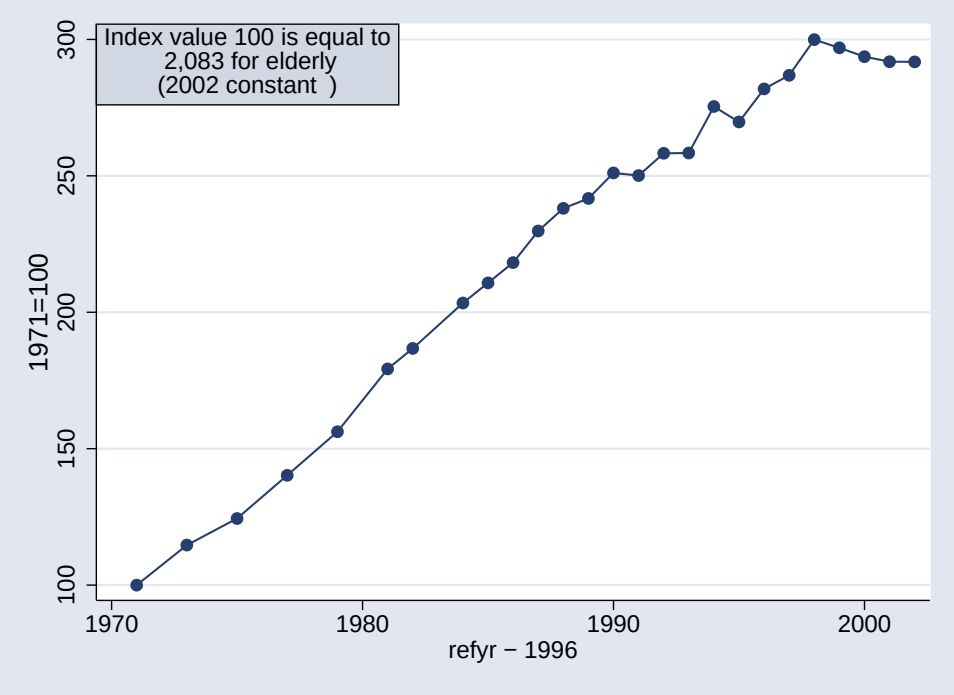


**Figure 7: After-tax income**

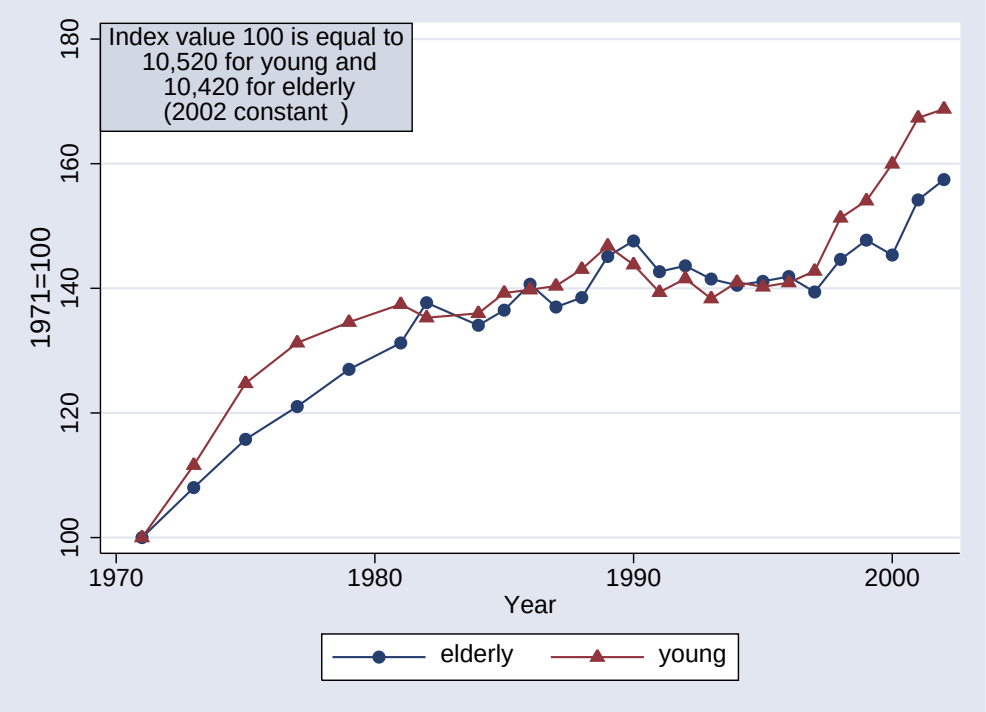

**Figure 8: Relative income poverty**

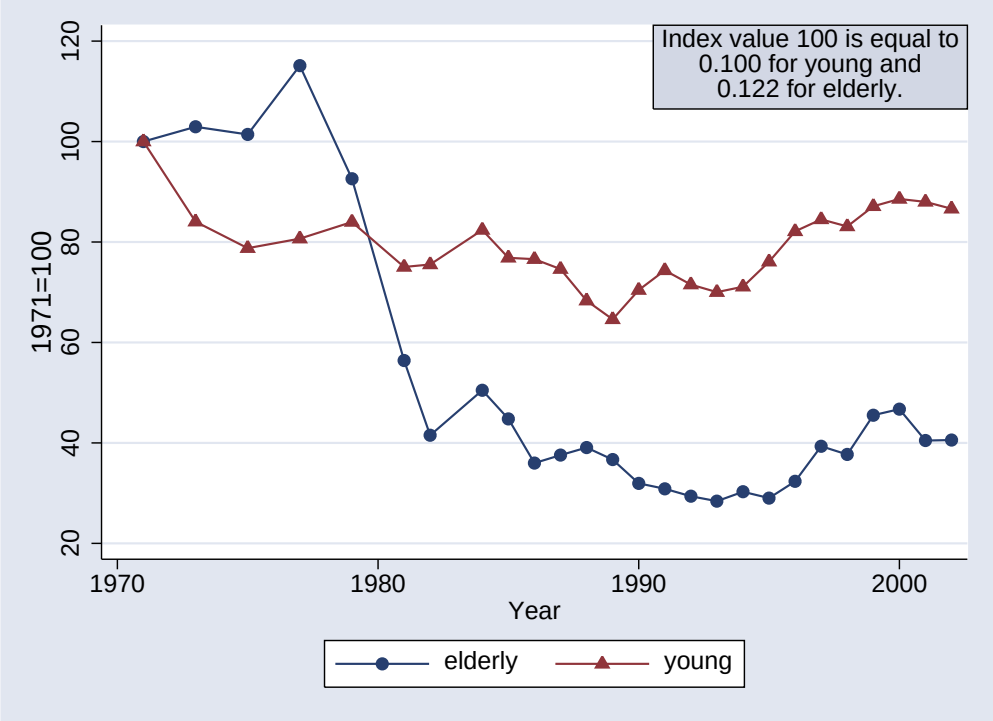


**Figure 9: Absolute income poverty**

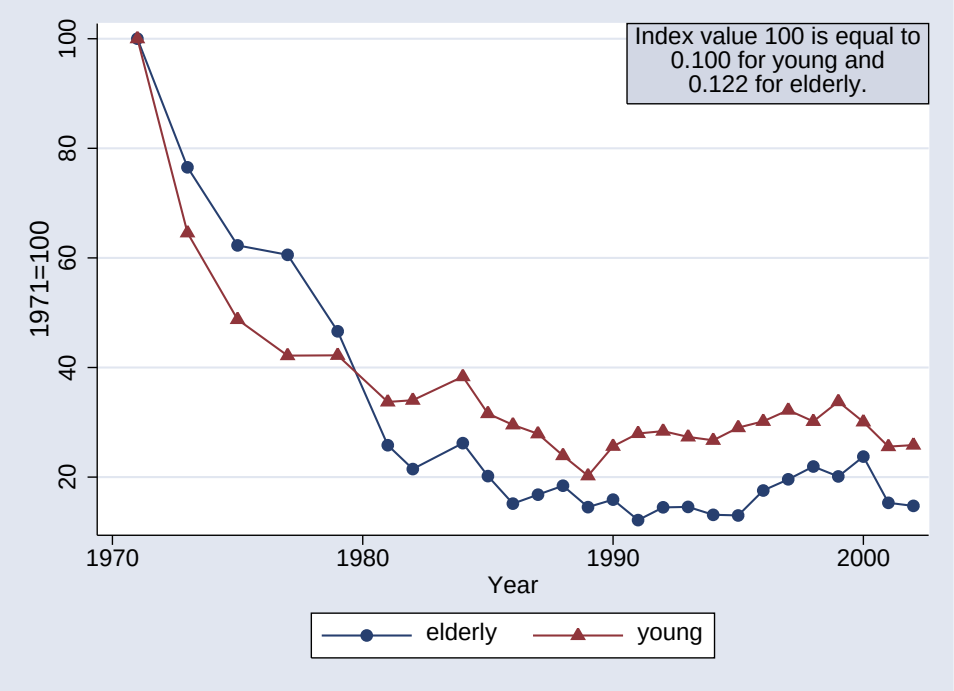

**Figure 10: Percentiles of Income**

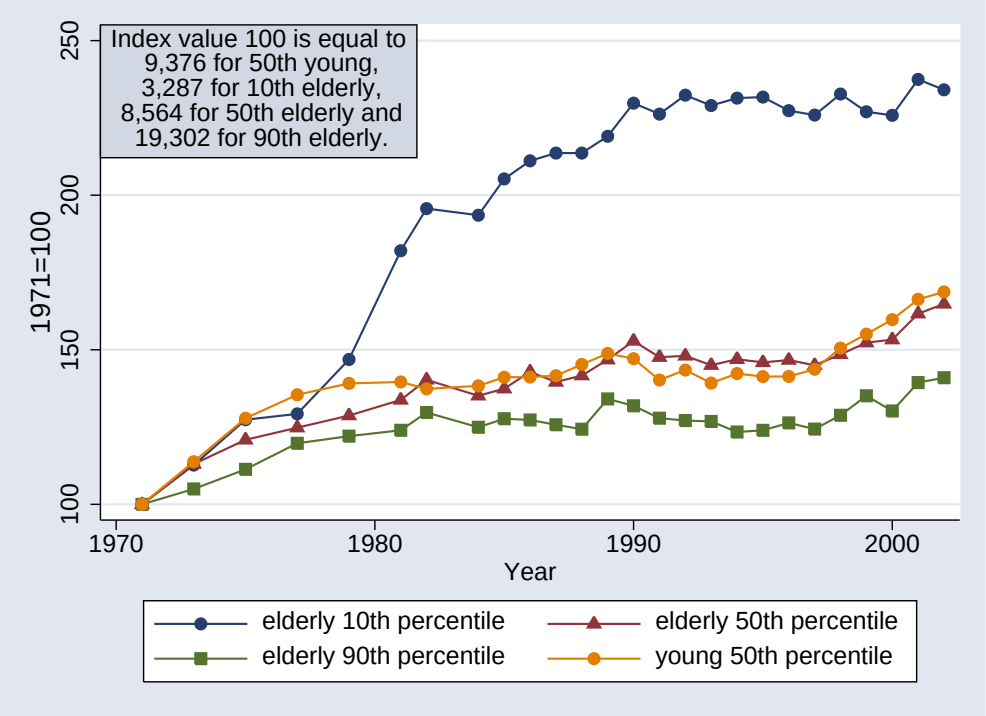


**Figure 11: Total Consumption**

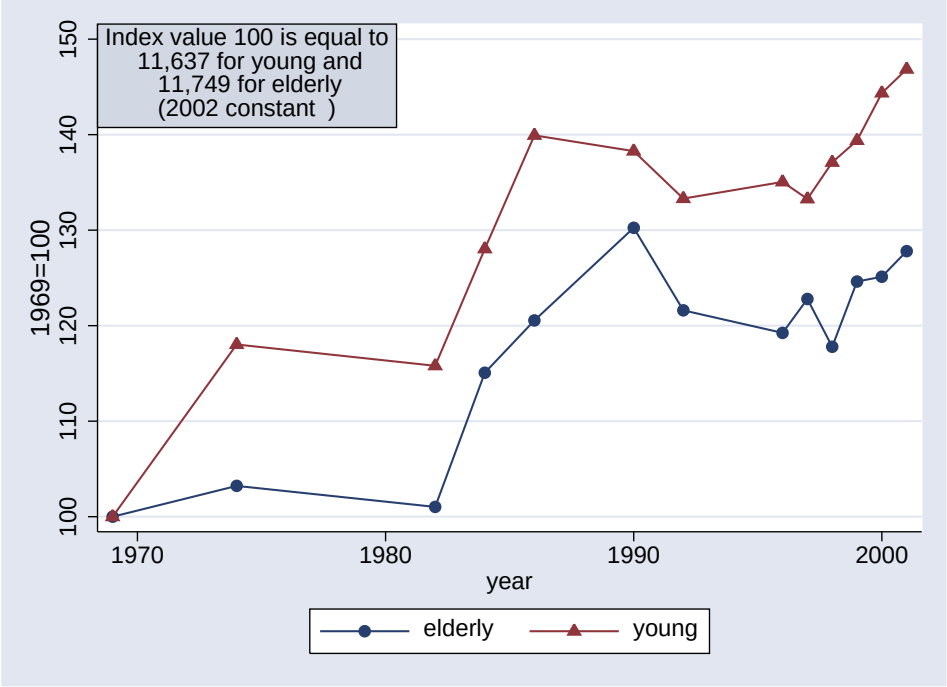

**Figure 12: Relative consumption poverty**

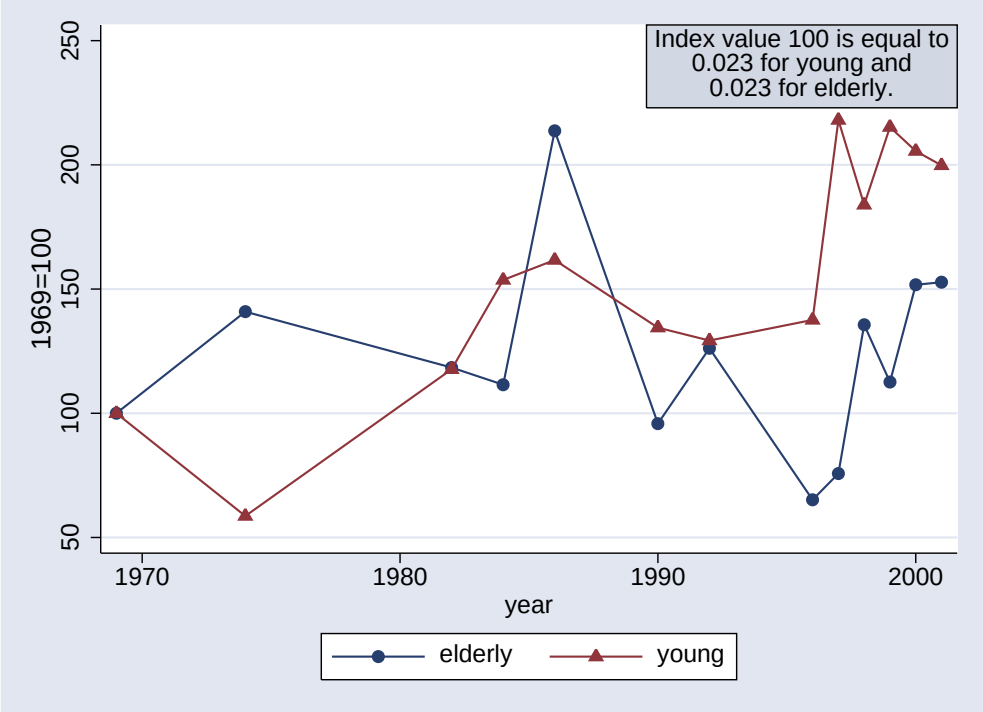


**Figure 13: Absolute consumption poverty**

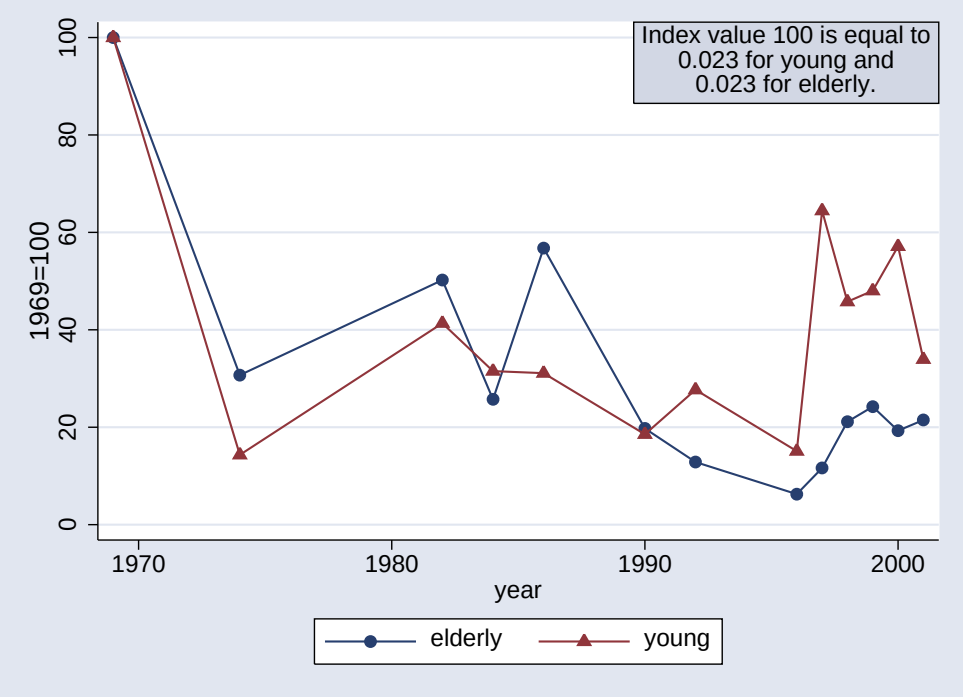

**Figure 14: Percentiles of Consumption**

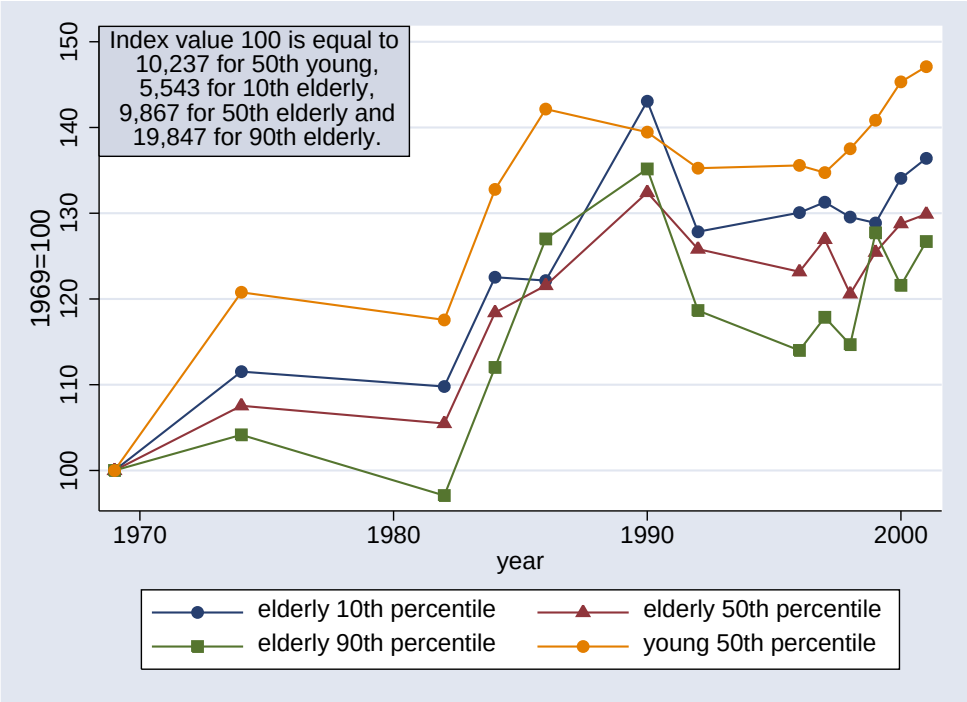


**Figure 15: Very Happy**

**Figure 16: Very Unhappy or Unhappy**

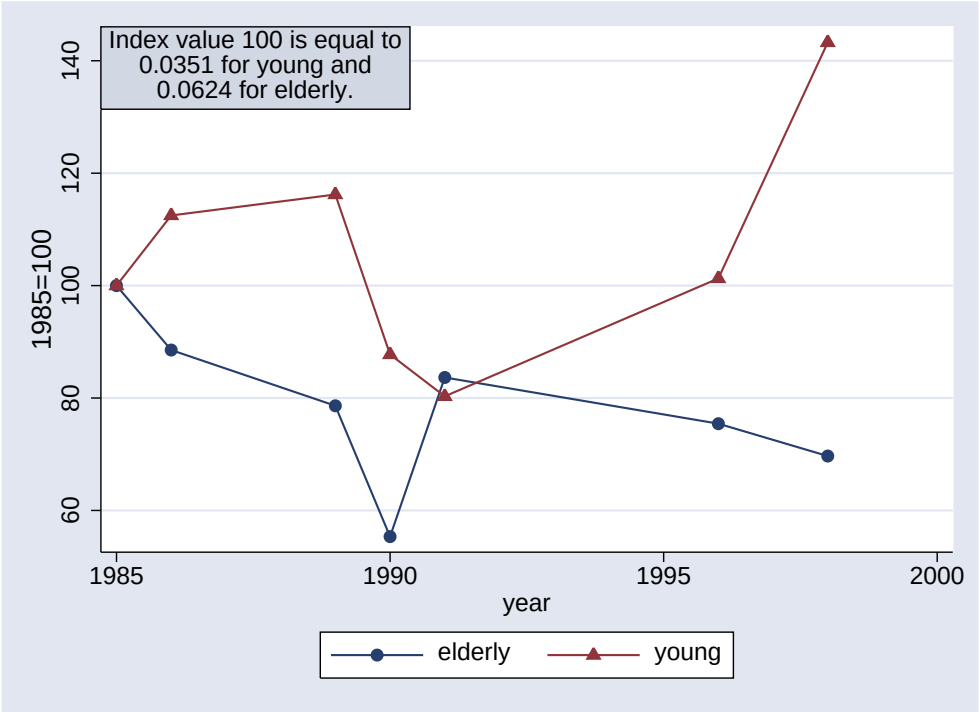


**Figure 17: Benefits vs. Income Scatterplot**

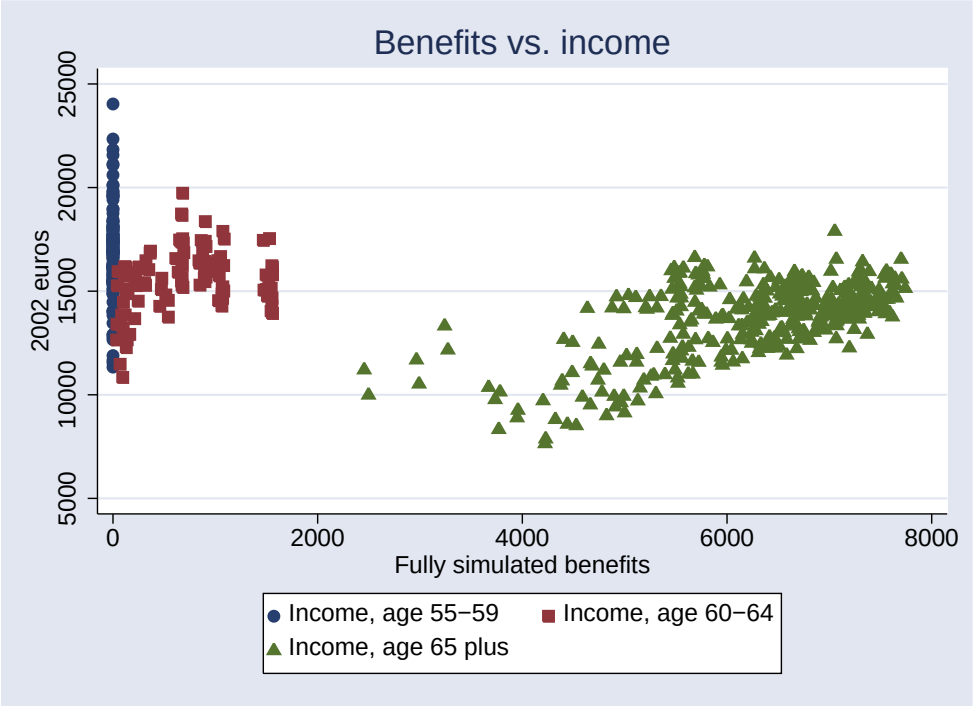

**Figure 18: Residuals vs. Benefits**

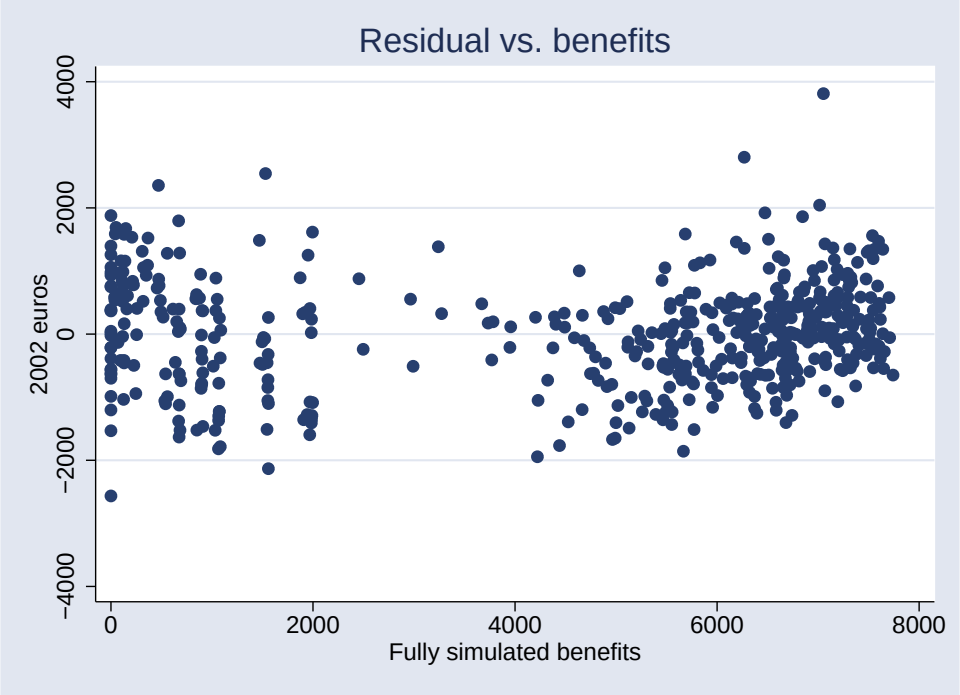


**Figure 19: Relative Income and Consumption Poverty by age in 2000**

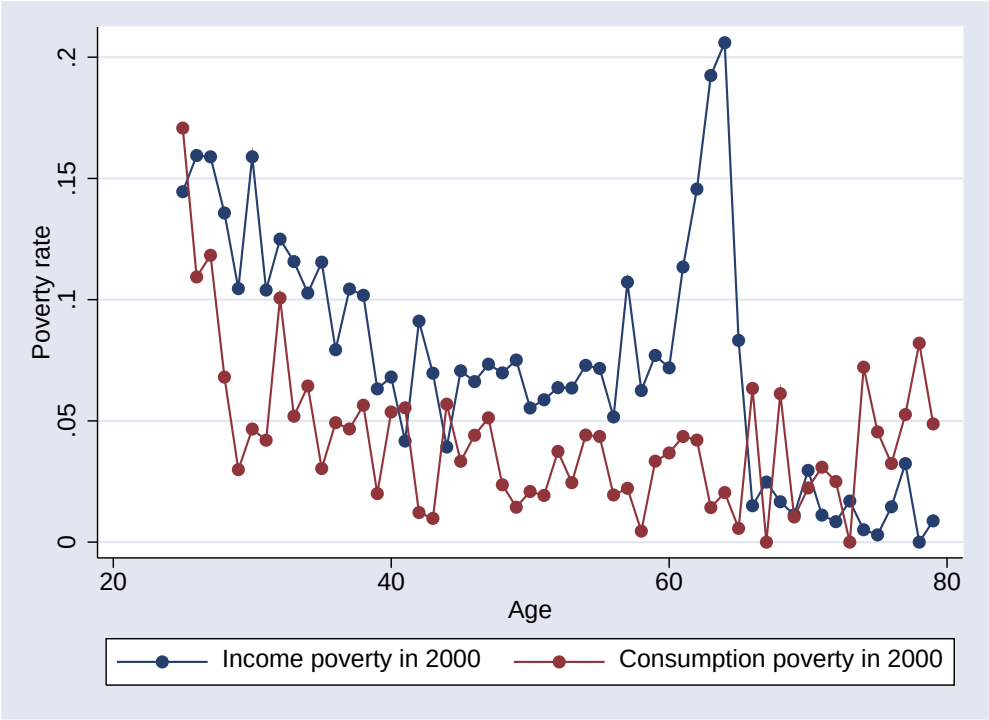

## Table 1: Data Description

| Measure | Data Source | Years Available | Ages Availab | Elderly Sample | Variable Description |
|---|---|---|---|---|---|
| Mean IS Income | SCF/SLID | 1971-2002 | 25-79 | 301,538 | Mean family IS income |
| Mean IS income below poverty line | SCF/SLID | 1971-2002 | 25-79 | 301,538 | Mean IS income among families below poverty line |
| IS income of 5th-15th percentile earners | SCF/SLID | 1971-2002 | 25-79 | 301,538 | Mean IS income for families between 5th and 15th percentile |
| IS income of 45th-55th percentile earners | SCF/SLID | 1971-2002 | 25-79 | 301,538 | Mean IS income for families between 45th and 55th percentile |
| IS income of 85th-95th percentile earners | SCF/SLID | 1971-2002 | 25-79 | 301,538 | Mean IS income for families between 85th and 95th percentile |
| Mean after-tax income | SCF/SLID | 1971-2002 | 25-79 | 301,538 | Mean total family income after tax |
| Relative income poverty | SCF/SLID | 1971-2002 | 25-79 | 301,538 | Equals one if under 40% of working family median income |
| Absolute income poverty | SCF/SLID | 1971-2002 | 25-79 | 301,538 | Equals one if under 40% of 1971 working family median income |
| 10th percentile income | SCF/SLID | 1971-2002 | 25-79 | 301,538 | 10th percentile of total after-tax family income |
| 50th percentile income | SCF/SLID | 1971-2002 | 25-79 | 301,538 | 50th percentile of total after-tax family income |
| 90th percentile income | SCF/SLID | 1971-2002 | 25-79 | 301,538 | 90th percentile of total after-tax family income |
| Mean consumption | FAMEX/SHS | 1969-2001 | 25-79 | 22,840 | Mean of total family expenditure |
| Relative consumption poverty | FAMEX/SHS | 1969-2001 | 25-79 | 22,840 | Equals one if under 40% of working family median consumption |
| Absolute consumption poverty | FAMEX/SHS | 1969-2001 | 25-79 | 22,840 | Equals one if under 40% of 1969 working family median consumption |
| 10th percentile consumption | FAMEX/SHS | 1969-2001 | 25-79 | 22,840 | 10th percentile of total family consumption |
| 50th percentile consumption | FAMEX/SHS | 1969-2001 | 25-79 | 22,840 | 50th percentile of total family consumption |
| 90th percentile consumption | FAMEX/SHS | 1969-2001 | 25-79 | 22,840 | 90th percentile of total family consumption |
| Very happy | GSS | 1985-1998 | 25-74 | 39,247 | Equals one if responded very happy |
| Unhappy or very unhappy | GSS | 1985-1998 | 25-74 | 39,247 | Equals one if responded unhappy or very unhappy |

Notes: SCF is Survey of Consumer Finances. SLID is Survey of Labour and Income Dynamics. FAMEX is Survey of Family Expenditures. SHS is survey of Household Spending. GSS is General Social Survey.

**Table 2: Income regression results**

| | | | First Stage | | Reduced Form | | IV | |
|---|---|---|---|---|---|---|---|---|
| | Mean | Obs. | Partial | Full | Partial | Full | Partial | Full |
| Mean IS income | 5135 | 500 | 0.759 | 0.830 | | | | |
| | | | (0.048) | (0.041) | | | | |
| Mean IS income, below poverty line | 1926 | 500 | -0.213 | -0.305 | | | | |
| | | | (0.162) | (0.182) | | | | |
| IS income of 5th-15th percentile earners | 2817 | 500 | 0.721 | 0.995 | | | | |
| | | | (0.140) | (0.164) | | | | |
| IS income of 45th-55th percentile earners | 5028 | 500 | 1.085 | 1.254 | | | | |
| | | | (0.118) | (0.115) | | | | |
| IS income of 85th-95th percentile earners | 7627 | 500 | 0.401 | 0.264 | | | | |
| | | | (0.041) | (0.043) | | | | |
| Mean income | 14313 | 500 | | | 0.415 | 0.890 | 0.546 | 1.073 |
| | | | | | (0.141) | (0.159) | (0.183) | (0.198) |
| Relative income poverty | 0.054 | 500 | | | -0.020 | -0.026 | -0.027 | -0.026 |
| | | | | | (0.004) | (0.004) | (0.006) | (0.005) |
| Absolute income poverty | 0.028 | 500 | | | -0.004 | -0.005 | -0.006 | -0.005 |
| | | | | | (0.004) | (0.004) | (0.006) | (0.005) |
| 10th percentile income | 6903 | 500 | | | 0.682 | 0.934 | 0.945 | 0.938 |
| | | | | | (0.134) | (0.164) | (0.106) | (0.083) |
| 50th percentile income | 12070 | 500 | | | 0.404 | 0.946 | 0.330 | 0.728 |
| | | | | | (0.138) | (0.151) | (0.110) | (0.133) |
| 90th percentile income | 24130 | 500 | | | 0.352 | 0.660 | 0.878 | 2.502 |
| | | | | | (0.233) | (0.230) | (0.608) | (1.048) |

Notes: Coefficients reported are for the variables listed on the left. Details on the specifications are in the text.

**Table 3: Consumption and Happiness Results**

| | | | Reduced Form | | IV | |
|---|---|---|---|---|---|---|
| | Mean | Obs. | Partial | Full | Partial | Full |
| Mean consumption | 13799 | 226 | 0.406 | 0.502 | 0.660 | 0.797 |
| | | | (0.188) | (0.191) | (0.333) | (0.329) |
| Relative consumption poverty | 0.029 | 226 | -0.004 | -0.005 | -0.008 | -0.009 |
| | | | (0.003) | (0.003) | (0.008) | (0.007) |
| Absolute consumption poverty | 0.007 | 226 | 0.000 | -0.001 | 0.000 | -0.001 |
| | | | (0.002) | (0.002) | (0.005) | (0.004) |
| 10th percentile consumption | 6963 | 226 | 0.180 | 0.218 | 0.385 | 0.383 |
| | | | (0.090) | (0.099) | (0.246) | (0.217) |
| 50th percentile consumption | 11852 | 226 | 0.170 | 0.311 | 0.150 | 0.275 |
| | | | (0.172) | (0.180) | (0.155) | (0.162) |
| 90th percentile consumption | 23030 | 226 | 0.542 | 0.689 | 1.177 | 2.484 |
| | | | (0.595) | (0.596) | (1.343) | (2.379) |
| Very happy | 0.610 | 28 | 0.023 | 0.008 | 0.122 | 0.026 |
| | | | (0.037) | (0.049) | (0.206) | (0.161) |
| Unhappy or very unhappy | 0.057 | 28 | -0.017 | -0.022 | -0.088 | -0.073 |
| | | | (0.008) | (0.011) | (0.064) | (0.054) |

Notes: Coefficients reported are for the variables listed on the left. Details on the specifications are in the text.